NSX Administration Guide

Modified on 28 MARCH 2023
VMware NSX 4.1



vmware®

You can find the most up-to-date technical documentation on the VMware website at:

https://docs.vmware.com/

VMware, Inc.
3401 Hillview Ave.
Palo Alto, CA 94304
www.vmware.com

Copyright © 2017-2023 VMware, Inc. All rights reserved. Copyright and trademark information.

Contents

About Administering VMware NSX 17

1 NSX Manager 18
- Security of NSX Manager 21
- License Enforcement in NSX Manager 22
- View Monitoring Dashboards 24

2 Tier-0 Gateways 27
- Add a Tier-0 Gateway 28
- Create an IP Prefix List 34
- Create a Community List 35
- Configure a Static Route 36
- Create a Route Map 37
- Using Regular Expressions to Match Community Lists When Adding Route Maps 39
- Configure BGP 40
  - BGP Autonomous System Number per Tier-0 VRF Gateway and BGP neighbor 48
- Tier-0 VRF Gateways 51
  - Deploy VRF-Lite with BGP 52
  - VRF Route Leaking 56
  - Inter-VRF Routing 59
- Configure OSPF 65
- Configure BFD 68
- Configure Multicast 69
- Configure IPv6 Layer 3 Forwarding 69
- Create SLAAC and DAD Profiles for IPv6 Address Assignment 70
- State Synchronization of Tier-0 Gateways 72
- Changing the HA Mode of a Tier-0 Gateway 72
- Configure the ARP Limit of a Tier-0 or Tier-1 Gateway or Logical Router 73
- Stateful Services on Tier-0 and Tier-1 73
  - Key Concepts Stateful Services 74
  - Supported Topologies 77
  - Configure Failure Domains 79
  - Configure Stateful Services on Tier-0 and Tier-1 Gateways 81
  - Understanding Traffic Flows 84

3 Tier-1 Gateway 87
- Add a Tier-1 Gateway 87
- State Synchronization of Tier-1 Gateways 90

4 Segments 92

Segment Profiles 92

Understanding QoS Segment Profile 94

Understanding IP Discovery Segment Profile 96

Understanding SpoofGuard Segment Profile 100

Understanding Segment Security Segment Profile 101

Understanding MAC Discovery Segment Profile 103

Add a Segment 104

Edge Bridging: Extending Overlay Segments to VLAN 108

Configure an Edge VM for Bridging 110

Create an Edge Bridge Profile 112

Extend an Overlay Segment to a VLAN or a Range of VLANs 115

Add a Metadata Proxy Server 116

Distributed Port Groups 117

5 DHCP 119

Configure DHCP Service 122

DHCP Configuration Settings: Reference 122

Configure Segment DHCP Server on a Segment 131

Configure Gateway DHCP Server on a Segment 133

Configure DHCP Relay on a Segment 135

Attach a DHCP Profile to a Tier-0 or Tier-1 Gateway 138

View Gateway DHCP Statistics 138

View Segment DHCP Statistics 140

Scenarios: Selection of Edge Cluster for DHCP Service 140

Scenarios: Impact of Changing Segment Connectivity on DHCP 145

6 Host Switches 148

Managing NSX on a vSphere Distributed Switch 148

Configuring a vSphere Distributed Switch 149

Managing NSX Distributed Virtual Port Groups 151

NSX Cluster Prepared with VDS 152

APIs to Configure vSphere Distributed Switch on NSX 153

Feature Support in a vSphere Distributed Switch Enabled to Support NSX 156

License for vSphere Distributed Switch 159

Enhanced Datapath 159

Automatically Assign ENS Logical Cores 159

Configure Guest Inter-VLAN Routing 160

Receive Side Scaling 163

Configure VTEP HA 169

VTEP Maintenance Mode 173

7 Virtual Private Network (VPN) 175
Understanding IPSec VPN 176
Using Policy-Based IPSec VPN 176
Using Route-Based IPSec VPN 178
Understanding Layer 2 VPN 179
Enable and Disable L2 VPN Path MTU Discovery 181
Adding VPN Services 182
Add an IPSec VPN Service 183
Add an L2 VPN Service 185
Adding IPSec VPN Sessions 188
Using Certificate-Based Authentication for IPSec VPN Sessions 188
Add a Policy-Based IPSec Session 189
Add a Route-Based IPSec Session 193
About Supported Compliance Suites 197
Adding L2 VPN Sessions 198
Add an L2 VPN Server Session 198
Add an L2 VPN Client Session 201
Download the Remote Side L2 VPN Configuration File 202
Add Local Endpoints 204
Adding Profiles 206
Add IKE Profiles 206
Add IPSec Profiles 209
Add DPD Profiles 211
Add an Autonomous Edge as an L2 VPN Client 212
Check the Realized State of an IPSec VPN Session 215
Understanding TCP MSS Clamping 218
Troubleshooting VPN Problems 219
Monitor and Troubleshoot VPN Sessions 219
Alarms When an IPsec VPN Session or Tunnel Is Down 219

8 Network Address Translation (NAT) 222
Configure NAT/DNAT/No SNAT/No DNAT/Reflexive NAT 224
Configure NAT64 227
NAT and Gateway Firewall 229

9 NSX Advanced Load Balancer (Avi) 231

10 Load Balancer 232
Key Load Balancer Concepts 233
Scaling Load Balancer Resources 233
Supported Load Balancer Features 234

Load Balancer Topologies 235
Setting Up Load Balancer Components 237
Add Load Balancers 237
Add an Active Monitor 239
Add a Passive Monitor 242
Add a Server Pool 244
Setting Up Virtual Server Components 247
Groups Created for Server Pools and Virtual Servers 279

11 Distributed Load Balancer 280
Understanding Traffic Flow with a Distributed Load Balancer 282
Create and Attach a Distributed Load Balancer Instance 283
Create a Server Pool for Distributed Load Balancer 284
Create a Virtual Server with a Fast TCP or UDP Profile 286
Verifying Distributed Load Balancer Configuration on ESXi Hosts 287
Distributed Load Balancer Statistics and Diagnostics 289
Distributed Load Balancer Operational Status 291
Run Traceflow on Distributed Load Balancer 294
Supported Features 295

12 Ethernet VPN (EVPN) 296
Overview of BGP EVPN 296
EVPN Support in NSX 299
EVPN Inline Mode 301
EVPN Inline Mode Configuration Workflow 302
EVPN Route Server Mode 307
EVPN Route Server Mode Configuration Workflow 308

13 Forwarding Policies 317
Add or Edit Forwarding Policies 318

14 IP Address Management (IPAM) 320
Add a DNS Zone 320
Add a DNS Forwarder Service 321
Add an IP Address Pool 322
Add an IP Address Block 323

15 Networking Settings 324
Configuring Multicast 324
Create an IGMP Profile 326
Create a PIM Profile 327

About IGMP Join 327
Add an EVPN/VXLAN VNI Pool 328
Configure Global Gateway Settings 328
Add a Gateway QoS Profile 329
Add a BFD Profile 330
Add a DHCP Profile 330
Add a DHCP Server Profile 331
Add a DHCP Relay Profile 333

16 Security 335
Firewall Rule Enforcement 336
Security Overview 336
NSX Guest Introspection Platform 342
NSX Guest Introspection Platform Architecture 342
NSX Guest Introspection Platform Use Cases 344
Installing Host Components 345
Installing Guest Components 345
Supported File Systems for Guest VMs 355
Logging and Troubleshooting Host Components 356
Logging and Troubleshooting Guest Components 361
Collect Environment and Workload Details 367
Supported Software 368
Security Monitoring 368
Using Aria Operations for Logs for Unified Security Logs 368
Monitoring Security Statistics 370
Security Terminology 372
Identity Firewall 372
Identity Firewall Workflow 373
Layer 7 Context Profile 376
Layer 7 Firewall Rule Workflow 377
Distributed Firewall 378
FQDN Filtering 378
Firewall Drafts 380
Configure Malicious IP Feeds 382
Malicious IPs Filtering and Analysis Dashboard 383
Add a Distributed Firewall 384
Distributed Firewall Packet Logs 388
Manage a Firewall Exclusion List 392
Extending Security Policies to Physical Workloads 393
Shared Address Sets 400
Export or Import a Firewall Configuration 400

Gateway Firewall 401
Supported Gateway Firewall Features on NSX Edge 402
Gateway Firewall Settings 403
Add a Gateway Firewall Policy and Rule 403
TLS Inspection 406
URL Filtering 419
FQDN Analysis 420
Gateway Firewall Packet Logs 422
Distributed Security for vSphere Distributed Switch 424
Install Distributed Security for vSphere Distributed Switch 426
Endpoint Protection 427
Understand Endpoint Protection 427
Configure Endpoint Protection 429
Manage Endpoint Protection 441
East-West Network Security - Chaining Third-party Services 456
Key Concepts of East-West Network Protection 456
NSX Requirements for East-West Traffic 457
High-Level Tasks for East-West Network Security 459
Deploy a Service for East-West Traffic Introspection 459
Add Redirection Rules for East-West Traffic 461
Exclude Members from a Security Service 463
Get List of Service Paths 464
Uninstall an East-West Traffic Introspection Service 465
Upgrade East-West Service VM 466
North-South Network Security - Inserting Third-party Service 470
High-Level Tasks for North-South Network Security 470
Deploy a Service for North-South Traffic Introspection 470
Add Redirection Rules for North-South Traffic 473
Uninstall a North-South Traffic Introspection Service 474
Update Service Insertion Status 474
Upgrade North-South Service VM 475
Network Introspection Settings 482
Add a Service Segment 482
Add a Service Profile 483
Add a Service Chain 483
NSX IDS/IPS and NSX Malware Prevention 484
Getting Started with NSX IDS/IPS and NSX Malware Prevention 485
Offline Downloading and Uploading NSX Intrusion Detection Signatures 504
Adding Security Profiles 508
Using NSX IDS/IPS and NSX Malware Prevention on a Distributed Firewall 510
Using NSX IDS/IPS and NSX Malware Prevention on a Gateway Firewall 530

Distributed IDS/IPS Logs 537
Monitoring File Events 540
Monitoring IDS/IPS Events 553
Administering NSX Malware Prevention 557
Troubleshooting NSX Malware Prevention 558
NSX Network Detection and Response 569
Getting Started with NSX Network Detection and Response 569
Working with the NSX Network Detection and Response Application 574
Administering NSX Network Detection and Response 663
Troubleshooting NSX Network Detection and Response 664
Time-Based Firewall Policy 673
Troubleshooting Firewall 675
Monitor and Troubleshoot Firewall on NSX Manager 675
Troubleshooting Distributed Firewall on ESX Hosts 675
Troubleshooting Gateway Firewall 685
Check Rule Realization Status 689
Distributed Firewall Packet Logs 691
Bare Metal Server Security 695
General Security Settings 696
Private IP Ranges 696
Firewall General Settings 696
Identity Firewall Event Log Sources 703
URL Database 703

17 Inventory 705
Add a Service 706
Add a Group 707
Overview of Group Membership Criteria 711
Profiles 719
Context Profiles 719
L7 Access Profiles 721
Attribute Types 722
App IDs 722
FQDNs 726
Custom URLs 727
URL Categories 731
Containers 731
Public Cloud Services 733
Physical Servers 733
Tags 733
Add Tags to an Object 737

Add a Tag to Multiple Objects 738
Unassign Tags from an Object 739
Unassign a Tag from Multiple Objects 740
Resource Sharing 740
Share Resources from the Default View 743

18 Multisite and NSX Federation 746
NSX Multisite 747
Working with VMware Site Recovery Manager and Multisite Environments 761
NSX Federation 761
Overview of NSX Federation 762
Networking in NSX Federation 777
Security in NSX Federation 794
Role-Based Access Control in NSX Federation 810
Traceflow in Federation 811
Prevent Password Lockout on Local Manager Nodes 814
Backup and Restore in NSX Federation 815
Disaster Recovery for Global Manager 817
Working with Site Recovery Manager and Federation 819
Network Recovery for Local Managers 821

19 Multi-tenancy 823
Projects 823
Add a Project 831
Add User Role Assignments for a Project 832
Add Quota to a Project 838
Monitor Quota Status of a Project 840
Networking Within a Project 841
Groups in a Project 843
Distributed Firewall Policies in a Project 844
Monitoring and Troubleshooting within Projects 846
Features Available for Consumption Under Projects 847

20 System Monitoring 851
Monitor NSX Edge Nodes and Gateways 851
APIs to Fetch Time-Series Metrics 854
Dynamic Plugins 859
Working with Events and Alarms 869
View Alarm Information 869
View Alarm Definitions 870
Configuring Alarm Definition Settings 872

Managing Alarm States 873
Registering Notification Watchers 874
Using Aria Operations for Logs or Splunk for System Monitoring 876
Using Aria Operations Manager for System Monitoring 882
Using Aria Operations for Networks for System Monitoring 886

21 Network Monitoring 897
Add an IPFIX Collector 897
Add a Firewall IPFIX Profile 898
Add a Switch IPFIX Profile 898
IPFIX Monitoring on a vSphere Distributed Switch 900
Add a Port Mirroring Session 901
Port Mirroring on a vSphere Distributed Switch 904
Perform a Traceflow 905
Simple Network Management Protocol (SNMP) 908
Network Latency Statistics 909
Measure Network Latency Statistics 913
Export Network Latency Statistics 915
Monitoring Tools in Manager Mode 917
View Port Connection Information in Manager Mode 917
Traceflow 917
Monitor Port Mirroring Sessions in Manager Mode 921
Configure Filters for a Port Mirroring Session 924
Configure IPFIX in Manager Mode 925
Monitor a Logical Switch Port Activity in Manager Mode 936
Checking CPU Usage and Network Latency 937
Live Traffic Analysis 938
Create a Live Traffic Analysis Session 940

22 Authentication and Authorization 943
Managing Local User Accounts 944
Add a Local User 944
Activate a Local User in NSX Manager 946
Manage Local User Accounts 946
Manage Local User's Password or Name Using the CLI 950
Delete a Local User 951
Password Management 952
Resetting User Passwords 952
Resetting Expired Passwords 954
Resetting the Passwords of an Appliance 955
Authentication Policy Settings 956

Integration with VMware Identity Manager/Workspace ONE Access 961
Time Synchronization between NSX Manager, vIDM, and Related Components 961
Obtain the Certificate Thumbprint from a vIDM Host 962
Configure VMware Identity Manager/Workspace ONE Access Integration 963
Validate VMware Identity Manager™ Functionality 965
Integration with LDAP 967
LDAP Identity Source 968
Configuring Both vIDM and LDAP or Transitioning from vIDM to LDAP 970
NSX API Authentication Using a Session Cookie 970
Add a Role Assignment or Principal Identity 976
Role-Based Access Control 979
Create or Manage Custom Roles 992
Logging User Account Changes 994

23 Certificates 995
Types of Certificates 995
Certificates for NSX and NSX Federation 997
Create a Certificate Signing Request File 1001
Creating Self-signed Certificates 1003
Create a Self-Signed Certificate 1003
Import a Certificate for a CSR 1004
Importing and Replacing Certificates 1004
Import a Self-signed or CA-signed Certificate 1005
Import a CA Certificate 1005
Set Checks for Certificate Imports 1006
Replace Certificates 1007
Importing and Retrieving CRLs 1010
Import a Certificate Revocation List 1010
Configuring NSX Manager to Retrieve a Certificate Revocation List 1011
Import or Update a Trusted CA Bundle 1011
Storage of Public Certificates and Private Keys for Load Balancer or VPN service 1012
Alarm Notification for Certificate Expiration 1012

24 Integration of Kubernetes Clusters with Antrea CNI 1013
Architecture of Antrea Kubernetes Cluster Integration with NSX 1014
Registering an Antrea Kubernetes Cluster to NSX 1016
Prerequisites for Registering an Antrea Kubernetes Cluster to NSX 1017
Edit the Bootstrap Configuration File 1022
Submit the YAML Files to the Kubernetes API Server 1025
Viewing Inventory of an Antrea Kubernetes Cluster in NSX Manager 1026
View Details of an Antrea Kubernetes Cluster 1027

View Details of Namespaces in an Antrea Kubernetes Cluster 1029
Monitor Health Status of an Antrea Kubernetes Cluster 1030
Trace the Path of a Packet with Antrea Traceflow 1033
Workaround to Trace the Packet Flow Between an Antrea Kubernetes Cluster and an NSX Overlay Network 1036
Antrea Groups 1037
Add an Antrea Group 1041
Generic Groups with Kubernetes Member Types in Dynamic Membership Criteria 1044
Distributed Firewall Policies for Securing Traffic Within an Antrea Kubernetes Cluster 1053
Add a Distributed Firewall Policy to Secure Traffic Within an Antrea Kubernetes Cluster 1055
Example: Add a Distributed Firewall Policy to Secure Traffic Within an Antrea Kubernetes Cluster 1061
Firewall Policies for Securing Traffic Between Antrea Kubernetes Clusters and VMs in an NSX Network 1063
Deregister an Antrea Kubernetes Cluster from NSX 1070
Delete an Antrea Kubernetes Cluster from the NSX Inventory by Using the Command Line 1072
Troubleshooting Antrea to NSX Integration Issues 1073
Collect Support Bundles for an Antrea Kubernetes Cluster 1074
Antrea Kubernetes Cluster Status is Down 1078

25 Configuring NSX in Manager Mode 1081
Logical Switches in Manager Mode 1081
Understanding BUM Frame Replication Modes 1082
Create a Logical Switch in Manager Mode 1084
Connecting a VM to a Logical Switch in Manager Mode 1085
Create a Logical Switch Port In Manager Mode 1093
Test Layer 2 Connectivity in Manager Mode 1094
Create a VLAN Logical Switch for the NSX Edge Uplink in Manager Mode 1097
Switching Profiles for Logical Switches and Logical Ports 1099
Edge Bridging in Manager Mode: Extending Overlay Segments to VLAN 1118
Logical Routers in Manager Mode 1126
Tier-1 Logical Router 1126
Tier-0 Logical Router 1137
NAT in Manager Mode 1169
Network Address Translation 1169
Grouping Objects in Manager Mode 1182
Create an IP Set in Manager Mode 1182
Create an IP Pool in Manager Mode 1183
Create a MAC Set in Manager Mode 1184
Create an NSGroup in Manager Mode 1184
Configuring Services and Service Groups 1186
Manage Tags for a VM in Manager Mode 1187

DHCP in Manager Mode 1188
DHCP 1189
Metadata Proxies 1193
IP Address Management in Manager Mode 1196
Manage IP Blocks in Manager Mode 1196
Manage Subnets for IP Blocks in Manager Mode 1196
Load Balancing in Manager Mode 1197
Key Load Balancer Concepts 1198
Configuring Load Balancer Components 1199
Firewall in Manager Mode 1229
Add or Delete a Firewall Rule to a Logical Router in Manager Mode 1229
Configure Firewall for a Logical Switch Bridge Port in Manager Mode 1230
Firewall Sections and Firewall Rules 1231
About Firewall Rules 1235

26 Backing Up and Restoring NSX Manager or Global Manager 1243
Configure Backups 1244
Start or Schedule Backups 1247
Remove Old Backups 1248
Listing Available Backups 1249
Restore a Backup 1250

27 Operations and Management 1254
View the Usage and Capacity of Categories of Objects 1255
Configuring the Login Banner and UI 1258
Configure the Login Window with a User Agreement Banner 1258
Configure the User Interface Settings 1259
Configure a Node Profile 1260
Checking the Realized State of a Configuration Change 1263
View Network Topology 1267
Search for Objects 1268
Filter by Object Attributes 1270
Add a Compute Manager 1270
Replace Compute Manager 1275
Configuring Active Directory and Event Log Scraping 1277
Enable Windows Security Log Access for the Event Log Reader 1279
Add an LDAP Server 1279
Synchronize Active Directory 1280
Remove NSX Extension from VMware vCenter 1281
Managing the NSX Manager Cluster 1282
View the Configuration and Status of the NSX Manager Cluster 1282

Update API Service Configuration of the NSX Manager Cluster 1285
Shut Down and Power On the NSX Manager Cluster 1286
Reboot an NSX Manager 1286
Change the IP Address of an NSX Manager 1287
Resize an NSX Manager Node 1289
Replacing an NSX Edge Transport Node in an NSX Edge Cluster 1291
Replace an NSX Edge Transport Node Using the NSX Manager UI 1291
Replace an NSX Edge Transport Node Using the API 1293
Managing Resource Reservations for an Edge VM Appliance 1297
Tune Resource Reservations for an NSX Edge Appliance 1298
Replacing NSX Edge Hardware or Redeploying NSX Edge Nodes VM 1299
Replace NSX Edge Hardware 1300
Redeploy an NSX Edge VM Appliance 1301
Configure Advanced NSX Edge Parameters to Improve Performance 1309
Edit NSX Edge VM Restart Policy 1312
Adding and Removing an ESXi Host Transport Node to and from vCenter Servers 1313
Changing the Distributed Router Interfaces' MAC Address 1314
Configuring Appliances 1315
Configuring NTP on Appliances and Transport Nodes 1315
Add a License Key and Generate a License Usage Report 1316
License Types 1319
Standards Compliance 1322
View Compliance Status Report 1323
Compliance Status Report Codes 1323
Configure Global FIPS Compliance Mode for Load Balancer 1328
Common Criteria Compliance 1331
Collect Support Bundles 1332
Understanding Support Bundle File Paths 1333
Debugging NSX at Runtime 1338
Log Messages and Error Codes 1348
Configure Remote Logging 1352
Add Syslog Servers for NSX Nodes 1359
Log Message IDs 1360
Troubleshooting Syslog Issues 1361
Configure Serial Logging on an Appliance VM 1362
Firewall Audit Log Messages 1362
Customer Experience Improvement Program 1379
Edit the Customer Experience Improvement Program Configuration 1379
Find the SSH Fingerprint of a Remote Server 1379
Configuring an External Load Balancer 1380
Configure Proxy Settings 1383

Promote Manager Objects to Policy Objects 1383
Back up and restore NSX configured in VMware vCenter 1386
Configure an External Entropy Source 1386

28 Using NSX Cloud 1388
Cloud Service Manager: UI Walkthrough 1388
Clouds 1389
System 1391
Threat Detection using the NSX Cloud Quarantine Policy 1396
Quarantine Policy in the NSX Enforced Mode 1397
Quarantine Policy in the Native Cloud Enforced Mode 1402
User Managed List for VMs 1402
NSX Enforced Mode 1403
Supported Operating Systems for Workload VMs 1404
Onboarding VMs in the NSX Enforced Mode 1406
Managing VMs in the NSX Enforced Mode 1415
Native Cloud Enforced Mode 1417
Managing VMs in the Native Cloud Enforced Mode 1417
NSX Features Supported with NSX Cloud 1420
Group VMs using NSX and Public Cloud Tags 1421
Use Native-Cloud Services 1425
Service Insertion for your Workload VMs in the NSX Enforced Mode 1426
Enable NAT on NSX-managed VMs 1435
Enable Syslog Forwarding 1436
Automate VPN for Public Cloud Endpoints using APIs 1436
Set up VPN in the Native Cloud Enforced Mode 1438
Set up VPN in the NSX Enforced Mode 1442
Deploying NSX Management Components On Microsoft Azure 1444
Redeploying Manager Nodes on Cloud Native Azure 1445
Managing Backup and Restore of NSX Manager and CSM in Microsoft Azure 1446
Restore CSM from Microsoft Azure Recovery Services Vault 1447
Restore NSX Manager from Microsoft Azure Recovery Services Vault 1449
NSX Cloud FAQs and Troubleshooting 1454

About Administering VMware NSX

The *NSX Administration Guide* provides information about configuring and managing networking for VMware NSX® (Formerly known as NSX-T Data Center), including how to create logical switches and ports and how to set up networking for tiered logical routers, configure NAT, firewalls, SpoofGuard, grouping and DHCP. It also describes how to configure NSX Cloud.

Intended Audience

This information is intended for anyone who wants to configure NSX. The information is written for experienced Windows or Linux system administrators who are familiar with virtual machine technology, networking, and security operations.

VMware Technical Publications Glossary

VMware Technical Publications provides a glossary of terms that might be unfamiliar to you. For definitions of terms as they are used in VMware technical documentation, go to https://www.vmware.com/topics/glossary.

Related Documentation

You can find the VMware NSX® Intelligence™ documentation at https://docs.vmware.com/en/VMware-NSX-Intelligence/index.html.

NSX Manager

1

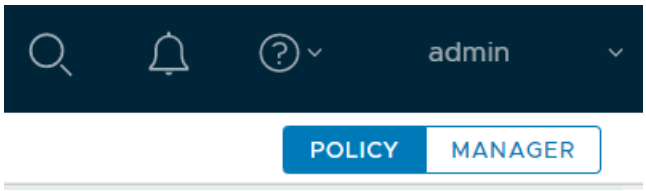
The NSX Manager provides a web-based user interface where you can manage your NSX environment. It also hosts the API server that processes API calls.

The NSX Manager interface provides two modes for configuring resources:

- Policy mode
- Manager mode

Accessing Policy Mode and Manager Mode

If present, you can use the **Policy** and **Manager** buttons to switch between the Policy and Manager modes. Switching modes controls which menus items are available to you.



- By default, if your environment contains only objects created through Policy mode, your user interface is in Policy mode and you do not see the **Policy** and **Manager** buttons.
- By default, if your environment contains any objects created through Manager mode, you see the **Policy** and **Manager** buttons in the top-right corner.

These defaults can be changed by modifying the user interface settings. See Configure the User Interface Settings for more information.

The same **System** tab is used in the Policy and Manager interfaces. If you modify Edge nodes, Edge clusters, or transport zones, it can take up to 5 minutes for those changes to be visible in Policy mode. You can synchronize immediately using `POST /policy/api/v1/infra/sites/default/enforcement-points/default?action=reload`.

When to Use Policy Mode or Manager Mode

Be consistent about which mode you use. There are a few reasons to use one mode over the other.

- If you are deploying a new NSX environment, using **Policy** mode to create and manage your environment is the best choice in most situations.
  - Some features are not available in Policy mode. If you need these features, use **Manager** mode for all configurations.
- If you plan to use NSX Federation, use **Policy** mode to create all objects. Global Manager supports only Policy mode.
- If you are upgrading from an earlier version of NSX and your configurations were created using the Advanced Networking & Security tab, use **Manager** mode.

  The menu items and configurations that were found under the Advanced Networking & Security tab are available in NSX 3.0 in **Manager** mode.

Important If you decide to use Policy mode, use it to create all objects. Do not use Manager mode to create objects.

Similarly, if you need to use Manager mode, use it to create all objects. Do not use Policy mode to create objects.

Table 1-1. When to Use Policy Mode or Manager Mode

Policy Mode	Manager Mode
Most new deployments should use Policy mode. NSX Federation supports only Policy mode. If you want to use NSX Federation, or might use it in future, use Policy mode.	Deployments which were created using the advanced interface, for example, upgrades from versions before Policy mode was available.
NSX Cloud deployments	Deployments which integrate with other plugins. For example, NSX Container Plug-in, Openstack, and other cloud management platforms.

Table 1-1. When to Use Policy Mode or Manager Mode (continued)

Policy Mode	Manager Mode
Networking features available in Policy mode only: ■ DNS Services and DNS Zones ■ VPN ■ Forwarding policies for NSX Cloud	

Security features available in Policy mode only:
■ Endpoint Protection
■ Network Introspection (East-West Service Insertion)
■ Context Profiles
  ■ L7 applications
  ■ FQDN
■ New Distributed Firewall and Gateway Firewall Layout
  ■ Categories
  ■ Auto service rules
  ■ Drafts
	Security features available in Manager mode only:
	■ Bridge Firewall

Names for Objects Created in Policy Mode and Manager Mode

The objects you create have different names depending on which interface was used to create them.

Table 1-2. Object Names

Objects Created Using Policy Mode	Objects Created Using Manager Mode
Segment	Logical switch
Tier-1 gateway	Tier-1 logical router
Tier-0 gateway	Tier-0 logical router
Group	NSGroup, IP Sets, MAC Sets
Security Policy	Firewall section
Gateway firewall	Edge firewall

Policy and Manager APIs

The NSX Manager provides two APIs: Policy and Manager.

- The Policy API contains URIs that begin with `/policy/api`.
- The Manager API contains URIs that begin with `/api`.

For more information about using the Policy API, see the NSX Policy API: Getting Started Guide.

Security

NSX Manager has the following security features:

- NSX Manager has a built-in user account called **admin**, which has access rights to all resources, but does not have rights to the operating system to install software. NSX upgrade files are the only files allowed for installation.
- NSX Manager supports session timeout and automatic user logout. NSX Manager does not support session lock. Initiating a session lock can be a function of the workstation operating system being used to access NSX Manager. Upon session termination or user logout, users are redirected to the login page.
- Authentication mechanisms implemented on NSX follow security best practices and are resistant to replay attacks. The secure practices are deployed systematically. For example, sessions IDs and tokens on NSX Manager for each session are unique and expire after the user logs out or after a period of inactivity. Also, every session has a time record and the session communications are encrypted to prevent session hijacking.

You can view and change the session timeout value with the following CLI commands:

- The command `get service http` displays a list of values including session timeout.
- To change the session timeout value, run the following commands:

```
set service http session-timeout <timeout-value-in-seconds>
restart service ui-service
```

This chapter includes the following topics:

- Security of NSX Manager
- License Enforcement in NSX Manager
- View Monitoring Dashboards

Security of NSX Manager

NSX Manager is a restricted system and has features designed to ensure the integrity of the system and to keep the system secure.

Details of the NSX Manager security features:

- NSX Manager supports session timeouts and user logoff. NSX Manager does not support session lock. To initiate a session lock use the workstation operating system being used to access NSX Manager.
- NSX Manager has one active local account out of the box, admin. You cannot delete or deactivate the admin account.

- There are three default inactive local user accounts: audit and two guest user accounts. In the Enterprise environment, the **guestuser1** and the **guestuser2** user accounts are available. In the NSX Cloud environment, the **cloud_admin** and the **cloud_audit** user accounts are available. You can delete, activate, or deactivate them. Starting in NSX 4.1, the admin in the Enterprise environment, as well as any user with the Enterprise Admin role, can add up to 10 more guest user accounts in addition to admin, audit, and two guest users.
- NSX Manager enforces approved authorizations for controlling the flow of management information within the network device based on information flow control policies.
- NSX Manager initiates session auditing upon startup.
- NSX Manager uses its internal system clock to generate time stamps for audit records.
- The NSX Manager user interface includes a user account, which has access rights to all resources, but does not have rights to the operating system to install software and hardware. NSX upgrade files are the only files allowed for installation. You cannot edit the rights of or delete this user.
- All passwords in the system (databases, configuration files, log files, and so on) get encrypted using a strong one-way hashing algorithm with a salt. During authentication, when the user enters the password it gets obfuscated. You can configure the password complexity using the UI, API and CLI. You can reset the node the authentication policy and the password complexity configuration back to their default system settings, if desired.
- FIPS compliance
  - NSX Manager uses FIPS 140-2 approved algorithms for authentication to a cryptographic module.
  - NSX Manager generates unique session identifiers using a FIPS 140-2 approved random number generator.
  - NSX Manager uses a FIPS 140-2 approved cryptographic algorithm to protect the confidentiality of remote maintenance and diagnostic sessions.
  - NSX Manager authenticates SNMP messages using a FIPS-validated Keyed-Hash Message Authentication Code (HMAC).
- NSX Manager recognizes only system-generated session identifiers and invalidates session identifiers upon administrator logout or other session termination.
- An audit log gets generated for events such as logon, logoff, and access to resources. Each audit log contains the timestamp, source, result, and a description of the event. For more information, see Log Messages and Error Codes.

License Enforcement in NSX Manager

License compliance is enforced when you try to access features in the NSX Manager user interface. The license enforcement is based on features and time, but not capacity. The capacity enforcement is honor-based.

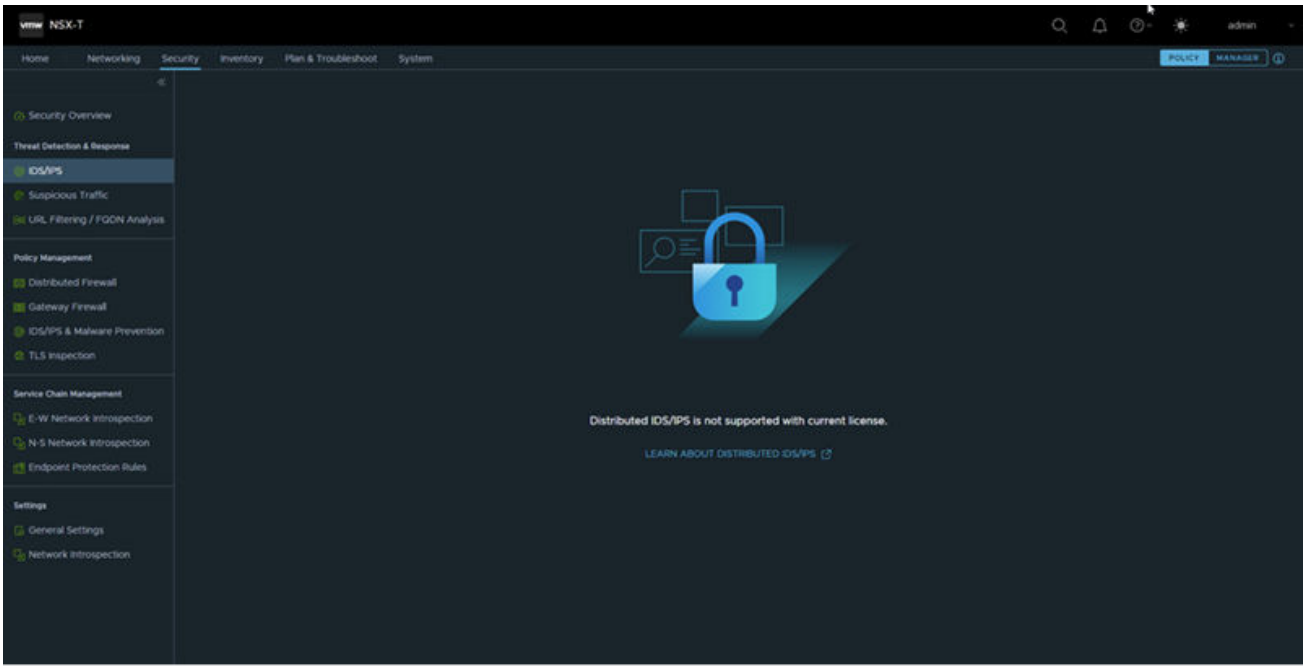
Beginning with NSX 3.1, the NSX license editions that you have assigned to your NSX deployment determine which features you can access in the Policy mode of the NSX Manager user interface. If you have multiple editions of licenses, NSX Manager uses the highest license edition that is applicable.

When the licenses are valid, the order of priority for the license editions is as follows.

Priority	License
1	NSX Data Center Enterprise Plus, NSX Data Center Evaluation
2	NSX Enterprise Plus per Processor (Limited Export), NSX Data Center Advanced, NSX for vSphere - Enterprise, NSX for vSphere - Advanced, NSX Data Center Advanced per Processor (for Limited Export)
3	NSX Data Center for Remote Office Branch Office (ROBO)
4	NSX Data Center Professional
5	NSX Data Center Standard and NSX for vSphere - Standard
6	NSX for vShield Endpoint

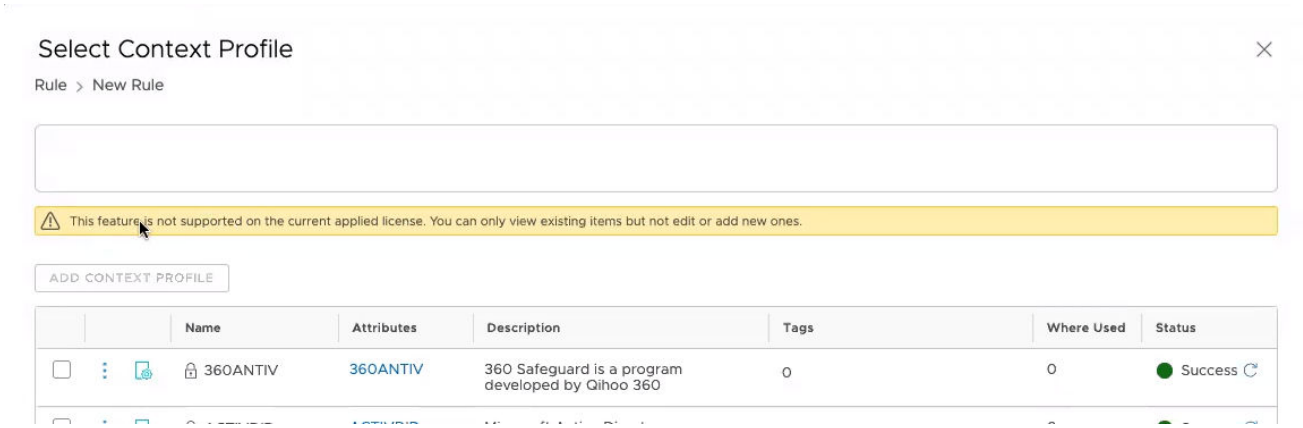
Note Add-On licenses verifies the add-on features, such as NSX Data Center Distributed Threat Prevention. For details, see License Types.

The assigned license is used to determine the list of features that you are allowed to use in the NSX Manager user interface. If you are a new user, you can only access those features that are available in the license edition that you have purchased. If you try to access a feature that is not valid for your current license, you see a message similar to what is shown in the following screenshot.



You can upgrade your current NSX deployment to NSX 3.1 or later regardless of the license that is in effect. Similarly, you can perform a backup or restore operation regardless of the license that is in effect. However, after a successful upgrade, backup, or restoration of a backup, the license enforcement is applied based on the current valid assigned license to your NSX deployment.

If an assigned license expires or becomes invalid, only the read and delete operations are allowed for objects that were configured before the license expired or before an upgrade, backup, or restoration process began. The edit, create, and new operations are disabled. A warning banner similar to what is shown in the following screenshot is displayed.



If NSX Manager has another valid license with a lower priority, alongside the expired higher priority license, then the features are enabled based on the priority of the valid license.

The list of supported features for the different VMware NSX Data Center license editions is available in the VMware NSX Data Center Datasheet at https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/products/nsx/vmware-nsx-datasheet.pdf. In that document, locate the VMware NSX Data Center Editions section and identify the license edition that is required for the features you want to use.

See Add a License Key and Generate a License Usage Report for more information on how to add a license key. You can also find information about the restrictions of the default license, NSX for vShield Endpoint, that is used when you install NSX Manager.

View Monitoring Dashboards

The NSX Manager interface provides numerous monitoring dashboards showing details regarding system status, networking and security, and compliance reporting. This information is displayed or accessible throughout the NSX Manager interface, but can be accessed together in the **Home > Monitoring Dashboards** page.

You can access the monitoring dashboards from the Home page of the NSX Manager interface. From the dashboards, you can click through and access the source pages from which the dashboard data is drawn.

Procedure

1 Log in as administrator to the NSX Manager interface.

2 Click **Home** if you are not already on the Home page.

3 Click **Monitoring Dashboards** and select the desired category of dashboards from the drop-down menu.

The page displays the dashboards in the selected categories. The dashboard graphics are color-coded, with color code key displayed directly above the dashboards.

4 To access a deeper level of detail, click the title of the dashboard, or one of the elements of the dashboard, if activated.

The following tables describe the default dashboards and their sources.

Table 1-3. System Dashboards

Dashboard	Sources	Description
System	**System > Appliances**	Shows the status of the NSX Manager cluster and NSX Advanced Load Balancer resource (CPU, memory, disk) consumption.
Fabric	**System > Fabric > Hosts** **System > Fabric > Nodes** **System > Fabric > Profiles** **System > Fabric > Transport Zones** **System > Fabric > Compute Managers** **System > Fabric > Settings**	Shows the status of the NSX fabric, including host and edge transport nodes, transport zones, and compute managers. You can also view profiles and set global fabric settings for tunnel endpoint, remote tunnel endpoint, and global MTU consistency checks if you select the Fabric title to access these additional tasks.
Backups	**System > Backup & Restore**	Shows the status of NSX backups, if configured. It is strongly recommended that you configure scheduled backups that are stored remotely to an SFTP site.
Endpoint Protection	**System > Service Deployments**	Shows the status of endpoint protection deployment.

Table 1-4. Networking & Security Dashboards in Policy Mode

Dashboard	Sources	Description
Security	**Inventory > Groups** **Security > Distributed Firewall**	Shows the status of groups and security policies. A group is a collection of workloads, segments, segment ports, IP addresses, MAC addresses, and so on where security policies, including East-West firewall rules, may be applied.
Gateways	**Networking > Tier-0 Gateways** **Networking > Tier-1 Gateways**	Shows the status of Tier-0 and Tier-1 gateways.
Segments	**Networking > Segments**	Shows the status of network segments.
Load Balancers	**Networking > Load Balancing**	Shows the status of the load balancer VMs.

Table 1-4. Networking & Security Dashboards in Policy Mode (continued)

Dashboard	Sources	Description
Virtual Services	**Networking > Network Services > Advanced Load Balancer/Load Balancing**	Shows the availability and scalability for virtual services applications by monitoring their health and distributing traffic.
VPNs	**Networking > VPN**	Shows the status of virtual private networks.

Table 1-5. Networking & Security Dashboards in Manager Mode

Dashboard	Sources	Description
Load Balancers	**Networking > Load Balancing**	Shows the status of the load balancer services, load balancer virtual servers, and load balancer server pools. A load balancer can host one or more virtual servers. A virtual server is bound to a server pool that includes members hosting applications.
Firewall	**Security > East West Security > Distributed Firewall** **Security > East West Security > Bridge Firewall** **Networking > Connectivity > Tier-0 Logical Routers** and **Networking > Connectivity > Tier-1 Logical Routers**	Indicates if the firewall is enabled, and shows the number of policies, rules, and exclusions list members. **Note** Each detailed item displayed in this dashboard is sourced from a specific sub-tab in the source page cited.
VPN	Not applicable.	Shows the status of virtual private networks and the number of IPSec and L2 VPN sessions open.
Switching	**Networking > Connectivity > Logical Switches**	Shows the status of logical switches and logical ports, including both VM and container ports.

Table 1-6. Compliance Report Dashboard

Column	Description
Non-Compliance Code	Displays the specific non-compliance code.
Description	Specific cause of non-compliance status.
Resource Name	The NSX resource (node, switch, and profile) in non-compliance.
Resource Type	Resource type of cause.
Affected Resources	Number of resources affected. Click the number value to view a list.

You can also add widget to configure custom monitoring dashboards using NSX REST APIs. See the latest version of the *NSX REST API Guide* at https://code.vmware.com for API details. See the Compliance Status Report Codes for more information about each compliance report code.

Tier-0 Gateways

2

A tier-0 gateway performs the functions of a tier-0 logical router. It processes traffic between the logical and physical networks.

NSX Cloud Note If using NSX Cloud, see NSX Features Supported with NSX Cloud for a list of auto-generated logical entities, supported features, and configurations required for NSX Cloud.

An Edge node can support only one tier-0 gateway or logical router. When you create a tier-0 gateway or logical router, make sure you do not create more tier-0 gateways or logical routers than the number of Edge nodes in the NSX Edge cluster.

Note When connecting tier-0 uplinks to multi-chassis port-channel topologies such as vPC (virtual PortChannel) or VSS (Virtual Switching System) from Cisco, or MLAG (Multi-Chassis Link Aggregation) from Arista, be sure to consult with the network provider to understand the limitations of the topology when it is being used for transit routing.

This chapter includes the following topics:

- Add a Tier-0 Gateway
- Create an IP Prefix List
- Create a Community List
- Configure a Static Route
- Create a Route Map
- Using Regular Expressions to Match Community Lists When Adding Route Maps
- Configure BGP
- Tier-0 VRF Gateways
- Configure OSPF
- Configure BFD
- Configure Multicast
- Configure IPv6 Layer 3 Forwarding
- Create SLAAC and DAD Profiles for IPv6 Address Assignment
- State Synchronization of Tier-0 Gateways

- Changing the HA Mode of a Tier-0 Gateway
- Configure the ARP Limit of a Tier-0 or Tier-1 Gateway or Logical Router
- Stateful Services on Tier-0 and Tier-1

Add a Tier-0 Gateway

A tier-0 gateway has downlink connections to tier-1 gateways and external connections to physical networks.

If you are adding a tier-0 gateway from Global Manager in NSX Federation, see Add a Tier-0 Gateway from Global Manager.

You can configure the HA (high availability) mode of a tier-0 gateway to be active-active or active-standby. The following services are only supported in active-standby mode:

- NAT
- Load balancing
- Stateful firewall
- VPN

Tier-0 and tier-1 gateways support the following addressing configurations for all interfaces (external interfaces, service interfaces and downlinks) in both single tier and multi-tiered topologies:

- IPv4 only
- IPv6 only
- Dual Stack - both IPv4 and IPv6

To use IPv6 or dual stack addressing, enable **IPv4 and IPv6** as the L3 Forwarding Mode in **Networking > Networking Settings > Global Networking Config**.

You can configure the tier-0 gateway to support EVPN (Ethernet VPN). For more information about configuring EVPN, see Chapter 12 Ethernet VPN (EVPN).

If you configure route redistribution for the tier-0 gateway, you can select from two groups of sources: tier-0 subnets and advertised tier-1 subnets. The sources in the tier-0 subnets group are:

Source Type	Description
Connected Interfaces and Segments	These include external interface subnets, service interface subnets and segment subnets connected to the tier-0 gateway.
Static Routes	Static routes that you have configured on the tier-0 gateway.
NAT IP	NAT IP addresses owned by the tier-0 gateway and discovered from NAT rules that are configured on the tier-0 gateway.
IPSec Local IP	Local IPSEC endpoint IP address for establishing VPN sessions.

Source Type	Description
DNS Forwarder IP	Listener IP for DNS queries from clients and also used as source IP used to forward DNS queries to upstream DNS server.
EVPN TEP IP	This is used to redistribute EVPN local endpoint subnets on the tier-0 gateway.

The sources in the advertised tier-1 subnets group are:

Source Type	Description
Connected Interfaces and Segments	These include segment subnets connected to the tier-1 gateway and service interface subnets configured on the tier-1 gateway.
Static Routes	Static routes that you have configured on the tier-1 gateway.
NAT IP	NAT IP addresses owned by the tier-1 gateway and discovered from NAT rules that are configured on the tier-1 gateway.
LB VIP	IP address of the load balancing virtual server.
LB SNAT IP	IP address or a range of IP addresses used for source NAT by the load balancer.
DNS Forwarder IP	Listener IP for DNS queries from clients and also used as source IP used to forward DNS queries to upstream DNS server.
IPSec Local Endpoint	IP address of the IPSec local endpoint.

Proxy ARP is automatically enabled on a tier-0 gateway when a NAT rule or a load balancer VIP uses an IP address from the subnet of the tier-0 gateway external interface. By enabling proxy-ARP, hosts on the overlay segments and hosts on a VLAN segment can exchange network traffic together without implementing any change in the physical networking fabric.

For a detailed example of a packet flow in a proxy ARP topology, see the *NSX Reference Design Guide* on the VMware Communities portal.

Before NSX 3.2, proxy ARP is supported on a tier-0 gateway in only an active-standby configuration, and it responds to ARP queries for the external and service interface IPs. Proxy ARP also responds to ARP queries for service IPs that are in an IP prefix list that is configured with the `Permit` action.

Starting in NSX 3.2, proxy ARP is also supported on a tier-0 gateway in an active-active configuration. However, all the Edge nodes in the active-active tier-0 configuration must have directly reachability to the network on which proxy ARP is required. In other words, you must configure the external interface and the service interface on all the Edge nodes that are participating in the tier-0 gateway for the proxy ARP to work.

Prerequisites

If you plan to configure multicast, see Configuring Multicast.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Select **Networking > Tier-0 Gateways**.
3. Click **Add Tier-0 Gateway**.

4 Enter a name for the gateway.

5 Select an HA (high availability) mode.

The default mode is active-active. In the active-active mode, traffic is load balanced across all members. In active-standby mode, all traffic is processed by an elected active member. If the active member fails, a new member is elected to be active.

6 If the HA mode is active-standby, select a failover mode.

Option	Description
Preemptive	If the preferred node fails and recovers, it will preempt its peer and become the active node. The peer will change its state to standby.
Non-preemptive	If the preferred node fails and recovers, it will check if its peer is the active node. If so, the preferred node will not preempt its peer and will be the standby node.

7 (Optional) Select an NSX Edge cluster.

8 (Optional) Click **Additional Settings**.

a In the **Internal Transit Subnet** field, enter a subnet.

This is the subnet used for communication between components within this gateway. The default is 169.254.0.0/24.

b In the **T0-T1 Transit Subnets** field, enter one or more subnets.

These subnets are used for communication between this gateway and all tier-1 gateways that are linked to it. After you create this gateway and link a tier-1 gateway to it, you will see the actual IP address assigned to the link on the tier-0 gateway side and on the tier-1 gateway side. The address is displayed in **Additional Settings > Router Links** on the tier-0 gateway page and the tier-1 gateway page. The default is 100.64.0.0/16.

After the tier-0 gateway is created, you can change the **T0-T1 Transit Subnets** by editing the gateway. Note that this will cause a brief disruption in traffic.

c In the **Forwarding Up Timer** field, enter a time.

Forwarding up timer defines the time in seconds that the router must wait before sending the up notification after the first BGP session is established. This timer (previously known as forwarding delay) minimizes downtime in case of fail-overs for active-active or active-standby configurations of logical routers on NSX Edge that use dynamic routing (BGP). It should be set to the number of seconds an external router (TOR) takes to advertise all the routes to this router after the first BGP/BFD session. The timer value should be directly proportional to the number of northbound dynamic routes that the router must learn. This timer should be set to 0 on single edge node setups.

9 Click **Route Distinguisher for VRF Gateways** to configure a route distinguisher admin address.

This is only needed for EVPN in Inline mode.

10 (Optional) Add one or more tags.

11 Click **Save**.

12 For IPv6, under **Additional Settings**, you can select or create an **ND Profile** and a **DAD Profile**.

These profiles are used to configure Stateless Address Autoconfiguration (SLAAC) and Duplicate Address Detection (DAD) for IPv6 addresses.

13 (Optional) Click **EVPN Settings** to configure EVPN.

a Select an EVPN mode.

The options are:

- **Inline** - In this mode, EVPN handles both data plane and control plane traffic.
- **Route Server** - Available only if this gateway's HA mode is active-active. In this mode, EVPN handles control plane traffic only.
- **No EVPN**

b If EVPN mode is **Inline**, select an EVPN/VXLAN VNI pool or create a new pool by clicking the menu icon (3 dots).

c If EVPN mode is **Route Server**, select an **EVPN Tenant** or create a new EVPN tenant by clicking the menu icon (3 dots).

d In the **EVPN Tunnel Endpoint** field click **Set** to add EVPN local tunnel endpoints.

For the tunnel endpoint, select an Edge node and specify an IP address.

Optionally, you can specify the MTU.

Note Ensure that the external interface has been configured on the NSX Edge node that you select for the EVPN tunnel endpoint.

14 To configure route redistribution, click **Route Redistribution** and **Set**.

Select one or more of the sources:

- Tier-0 subnets: **Static Routes**, **NAT IP**, **IPSec Local IP**, **DNS Forwarder IP**, **EVPN TEP IP**, **Connected Interfaces & Segments**.

  Under **Connected Interfaces & Segments**, you can select one or more of the following: **Service Interface Subnet**, **External Interface Subnet**, **Loopback Interface Subnet**, **Connected Segment**.

- Advertised tier-1 subnets: **DNS Forwarder IP**, **Static Routes**, **LB VIP**, **NAT IP**, **LB SNAT IP**, **IPSec Local Endpoint**, **Connected Interfaces & Segments**.

  Under **Connected Interfaces & Segments**, you can select **Service Interface Subnet** and/or **Connected Segment**.

15 To configure interfaces, click **Interfaces** and **Set**.

a Click **Add Interface**.

b Enter a name.

c Select a type.

If the HA mode is active-standby, the choices are **External**, **Service**, and **Loopback**. If the HA mode is active-active, the choices are **External** and **Loopback**.

d Enter an IP address in CIDR format.

e Select a segment.

f If the interface type is not **Service**, select an NSX Edge node.

g (Optional) If the interface type is not **Loopback**, enter an MTU value.

h (Optional) If the interface type is **External**, you can enable multicast by setting **PIM** (Protocol Independent Multicast) to **Enabled**.

You can also configure the following:

- **IGMP Join Local** - Enter one or more IP addresses. IGMP join is a debugging tool used to generate (*,g) join to Rendezvous Point (RP) and get traffic forwarded to the node where the join is issued. For more information, see About IGMP Join.
- **Hello Interval (seconds)** - Default is 30. The range is 1 - 180. This parameter specifies the time between Hello messages. After the **Hello Interval** is changed, it takes effect only after the currently scheduled PIM timer expires
- **Hold Time (seconds)** - The range is 1 - 630. Must be greater than **Hello Interval**. The default is 3.5 times **Hello Interval**. If a neighbor does not receive a Hello message from this gateway during this time interval, the neighbor will consider this gateway unreachable.

i (Optional) Add tags and select an ND profile.

j (Optional) If the interface type is **External**, for **URPF Mode**, you can select **Strict** or **None**.

URPF (Unicast Reverse Path Forwarding) is a security feature.

k (Optional) After you create an interface, you can download the aggregate of ARP proxies for the gateway by clicking the menu icon (three dots) for the interface and selecting **Download ARP Proxies**.

You can also download the ARP proxy for a specific interface by expanding a gateway and then expanding **Interfaces**. Click an interface and click the menu icon (three dots) and select **Download ARP Proxy**.

Note You cannot download the ARP proxy for loopback interfaces.

16 (Optional) If the HA mode is active-standby, click **Set** next to **HA VIP Configuration** to configure HA VIP.

With HA VIP configured, the tier-0 gateway is operational even if one external interface is down. The physical router interacts with the HA VIP only. HA VIP is intended to work with static routing and not with BGP.

a Click **Add HA VIP Configuration**.

b Enter an IP address and subnet mask.

The HA VIP subnet must be the same as the subnet of the interface that it is bound to.

c Select two interfaces from two different Edge nodes.

17 Click **Routing** to add IP prefix lists, community lists, static routes, and route maps.

18 Click **Multicast** to configure multicast routing.

19 Click **BGP** to configure BGP.

20 Click **OSPF** to configure OSPF.

This feature is available starting with NSX 3.1.1.

21 (Optional) To download the routing table or forwarding table, do the following:

a Click the menu icon (three dots) and select a download option.

b Enter values for **Transport Node**, **Network**, and **Source** as required.

c Click **Download** to save the .CSV file.

22 (Optional) To download the ARP table from a linked tier-1 gateway, do the following:

a From the **Linked Tier-1 Gateways** column, click the number.

b Click the menu icon (3 dots) and select **Download ARP Table**.

c Select an edge node.

d Click **Download** to save the .CSV file.

Results

The new gateway is added to the list. For any gateway, you can modify its configurations by clicking the menu icon (3 dots) and select **Edit**. For the following configurations, you do not need to click **Edit**. You only need to click the expand icon (right arrow) for the gateway, find the entity and click the number next to it. Note that the number must be non-zero. If it is zero, you must edit the gateway.

- In the **Interfaces** section: **External and Service Interfaces**.
- In the **Routing** section: **IP Prefix Lists**, **Static Routes**, **Static Route BFD Peer**, **Community Lists**, **Route Maps**.
- In the **BGP** section: **BGP Neighbors**.

If NSX Federation is configured, this feature of reconfiguring a gateway by clicking on an entity is applicable to gateways created by the Global Manager (GM) as well. Note that some entities in a GM-created gateway can be modified by the Local Manager, but others cannot. For example, **IP Prefix Lists** of a GM-created gateway cannot be modified by the Local Manager. Also, from the Local Manager, you can edit existing **External and Service Interfaces** of a GM-created gateway but you cannot add an interface.

Create an IP Prefix List

An IP prefix list contains single or multiple IP addresses that are assigned access permissions for route advertisement. The IP addresses in this list are processed sequentially. IP prefix lists are referenced through BGP neighbor filters or route maps with in or out direction.

For example, you can add the IP address 192.168.100.3/27 to the IP prefix list and deny the route from being redistributed to the northbound router. You can also append an IP address with less-than-or-equal-to (le) and greater-than-or-equal-to (ge) modifiers to grant or limit route redistribution. For example, 192.168.100.3/27 ge 24 le 30 modifiers match subnet masks greater than or equal to 24-bits and less than or equal to 30-bits in length.

Note The default action for a route is **Deny**. When you create a prefix list to deny or permit specific routes, be sure to create an IP prefix with no specific network address (select **Any** from the dropdown list) and the **Permit** action if you want to permit all other routes.

Prerequisites

Verify that you have a tier-0 gateway configured. See Create a Tier-0 Logical Router in Manager Mode.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Select **Networking > Tier-0 Gateways**.
3. To edit a tier-0 gateway, click the menu icon (three dots) and select **Edit**.
4. Click **Routing**.
5. Click **Set** next to **IP Prefix List**.
6. Click **Add IP Prefix List**.
7. Enter a name for the IP prefix list.
8. Click **Set** to add IP prefixes.

9 Click **Add Prefix**.

   a Enter an IP address in CIDR format.

     For example, 192.168.100.3/27.

   b (Optional) Set a range of IP address numbers in the **le** or **ge** modifiers.

     For example, set **le** to 30 and **ge** to 24.

   c Select **Deny** or **Permit** from the drop-down menu.

   d Click **Add**.

10 Repeat the previous step to specify additional prefixes.

11 Click **Save**.

Create a Community List

You can create BGP community lists so that you can configure route maps based on community lists.

Community lists are user-defined lists of community attribute values. These lists can be used for matching or manipulating the communities attribute in BGP update messages.

Both the BGP Communities attribute (RFC 1997) and the BGP Large Communities attribute (RFC 8092) are supported. The BGP Communities attribute is a 32-bit value split into two 16-bit values. The BGP Large Communities attribute has 3 components, each 4 octets in length.

In route maps we can match on or set the BGP Communities or Large Communities attribute. Using this feature, network operators can implement network policy based on the BGP communities attribute.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Networking > Tier-0 Gateways**.

3 To edit a tier-0 gateway, click the menu icon (three dots) and select **Edit**.

4 Click **Routing**.

5 Click **Set** next to **Community List**.

6 Click **Add Community List**.

7 Enter a name for the community list.

8 Specify a list of communities. For a regular community, use the aa:nn format, for example, 300:500. For a large community, use the format aa:bb:cc, for example, 11:22:33. Note that the list cannot have both regular communities and large communities. It must contain only regular communities, or only large communities.

   In addition, you can select one or more of the following regular communities. Note that they cannot be added if the list contains large communinities.

   - NO_EXPORT_SUBCONFED - Do not advertise to EBGP peers.
   - NO_ADVERTISE - Do not advertise to any peer.
   - NO_EXPORT - Do not advertise outside BGP confederation

9 Click **Save**.

Configure a Static Route

You can configure a static route on the tier-0 gateway to external networks. After you configure a static route, there is no need to advertise the route from tier-0 to tier-1, because tier-1 gateways automatically have a static default route towards their connected tier-0 gateway.

Recursive static routes are supported.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Networking > Tier-0 Gateways**.

3 To edit a tier-0 gateway, click the menu icon (three dots) and select **Edit**.

4 Click **Routing**.

5 Click **Set** next to **Static Routes**.

6 Click **Add Static Route**.

7 Enter a name and network address in CIDR format. Static routes based on IPv6 are supported. IPv6 prefixes can only have an IPv6 next hop.

8 Click **Set Next Hops** to add next-hop information.

9 Click **Add Next Hop**.

10 Enter an IP address or select **NULL**.

   If **NULL** is selected, the route is called a device route.

11 Specify the administrative distance.

12 Select a scope from the drop-down list. A scope can be an interface, a gateway, an IPSec session, or a segment.

13 Click **Add**.

What to do next

Check that the static route is configured properly. See Verify the Static Route on a Tier-0 Router.

Create a Route Map

A route map consists of a sequence of IP prefix lists, BGP path attributes, and an associated action. The router scans the sequence for an IP address match. If there is a match, the router performs the action and scans no further.

Route maps can be referenced at the BGP neighbor level and for route redistribution.

Prerequisites

- Verify that an IP prefix list or a community list is configured. See Create an IP Prefix List in Manager Mode or Create a Community List.
- For details about using regular expressions to define route-map match criteria for community lists, see Using Regular Expressions to Match Community Lists When Adding Route Maps.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Select **Networking > Tier-0 Gateways**.
3. To edit a tier-0 gateway, click the menu icon (three dots) and select **Edit**.
4. Click **Routing**.
5. Click **Set** next to **Route Maps**.
6. Click **Add Route Map**.
7. Enter a name and click **Set** to add match criteria.
8. Click **Add Match Criteria** to add one or more match criteria.

9 For each criterion, select **IP Prefix** or **Community List** and click **Set** to specify one or more match expressions.

a If you selected **Community List**, specify match expressions that define how to match members of community lists. For each community list, the following match options are available:

- **MATCH ANY** - perform the set action in the route map if any of the communities in the community list is matched.
- **MATCH ALL** - perform the set action in the route map if all the communities in the community list are matched regardless of the order.
- **MATCH EXACT** - perform the set action in the route map if all the communities in the community list are matched in the exact same order.
- **MATCH COMMUNITY REGEXP** - perform the set action in the route map if all the regular communities associated with the NRLI match the regular expression.
- **MATCH LARGE COMMUNITY REGEXP** - perform the set action in the route map if all the large communities associated with the NRLI match the regular expression.

You should use the match criterion MATCH_COMMUNITY_REGEX to match routes against standard communities, and use the match criterion MATCH_LARGE_COMMUNITY_REGEX to match routes against large communities. If you want to permit routes containing either the standard community or large community value, you must create two match criteria. If the match expressions are given in the same match criterion, only the routes containing both the standard and large communities will be permitted.

For any match criterion, the match expressions are applied in an AND operation, which means that all match expressions must be satisfied for a match to occur. If there are multiple match criteria, they are applied in an OR operation, which means that a match will occur if any one match criterion is satisfied.

10 Set BGP attributes.

BGP Attribute	Description
AS-path Prepend	Prepend a path with one or more AS (autonomous system) numbers to make the path longer and therefore less preferred.
MED	Multi-Exit Discriminator indicates to an external peer a preferred path to an AS.
Weight	Set a weight to influence path selection. The range is 0 - 65535.

BGP Attribute	Description
Community	Specify a list of communities. For a regular community use the aa:nn format, for example, 300:500. For a large community use the aa:bb:cc format, for example, 11:22:33. Or use the drop-down menu to select one of the following: ■ NO_EXPORT_SUBCONFED - Do not advertise to EBGP peers. ■ NO_ADVERTISE - Do not advertise to any peer. ■ NO_EXPORT - Do not advertise outside BGP confederation
Local Preference	Use this value to choose the outbound external BGP path. The path with the highest value is preferred.

11 In the Action column, select **Permit** or **Deny**.

You can permit or deny IP addresses matched by the IP prefix lists or community lists from being advertised.

12 Click **Save**.

Using Regular Expressions to Match Community Lists When Adding Route Maps

You can use regular expressions to define the route-map match criteria for community lists. BGP regular expressions are based on POSIX 1003.2 regular expressions.

The following expressions are a subset of the POSIX regular expressions.

Expression	Description
.*	Matches any single character.
*	Matches 0 or more occurrences of pattern.
+	Matches 1 or more occurrences of pattern.
?	Matches 0 or 1 occurrence of pattern.
^	Matches the beginning of the line.
$	Matches the end of the line.
_	This character has special meanings in BGP regular expressions. It matches to a space, comma, AS set delimiters { and } and AS confederation delimiters (and). It also matches to the beginning of the line and the end of the line. Therefore this character can be used for an AS value boundaries match. This character technically evaluates to (^|[,{}()]|$).

Here are some examples for using regular expressions in route maps:

Expression	Description
^101	Matches routes having community attribute that starts with 101.
^[0-9]+	Matches routes having community attribute that starts with a number between 0-9 and has one or more instances of such a number.
.*	Matches routes having any or no community attribute.

Expression	Description
.+	Matches routes having any community value.
^$	Matches routes having no/null community value.

Configure BGP

To enable access between your VMs and the outside world, you can configure an external or internal BGP (eBGP or iBGP) connection between a tier-0 gateway and a router in your physical infrastructure.

When configuring BGP, you must configure a local Autonomous System (AS) number for the tier-0 gateway. You must also configure the remote AS number. EBGP neighbors must be directly connected and in the same subnet as the tier-0 uplink. If they are not in the same subnet, BGP multi-hop should be used.

BGPv6 is supported for single hop and multihop. Redistribution, prefix list, and route maps are supported with IPv6 prefixes.

RFC-5549 enables BGPv6 sessions to exchange IPv4 routes with an IPv6 next hop. To minimize the number of BGP sessions and IPv4 addresses, you can exchange both IPv4 and IPv6 routes over a BGP session. Support for encoding and processing an IPv4 route with an IPv6 next hop is negotiated as part of the capability exchange in the BGP OPEN message. If both sides of a peering session support the capability, IPv4 routes are advertised with an IPv6 next hop. Multi-protocol BGP (MP-BGP) is used to advertise the Network Layer Reachability Information of a IPv4 address family using the next hop of an IPv6 address family.

A tier-0 gateway in active-active mode supports inter-SR (service router) iBGP. If gateway #1 is unable to communicate with a northbound physical router, traffic is re-routed to gateway #2 in the active-active cluster. If gateway #2 is able to communicate with the physical router, traffic between gateway #1 and the physical router will not be affected. A route learned by an Edge node from a northbound router will always be preferred to the same route learned over inter-SR iBGP. It is not possible to change this preference.

The implementation of ECMP on NSX Edge is based on the 5-tuple of the protocol number, source and destination address, and source and destination port.

The iBGP feature has the following capabilities and restrictions:

- Redistribution, prefix lists, and routes maps are supported.
- Route reflectors are not supported.
- BGP confederation is not supported.

NSX 4.1.0 introduces the support for:

- **BGP Autonomous System Number (ASN) per Tier-0 VRF Gateway and BGP neighbor**: You can configure a different BGP ASN per Tier-0 VRF Gateway and also per BGP neighbor. For more information, see BGP Autonomous System Number per Tier-0 VRF Gateway and BGP neighbor.

- **Inter-VRF Routing**: You can configure inter-VRF routing using easier workflows by importing and exporting routes between VRFs. For more information, see Inter-VRF Routing.
- **Autonomous-System-Wide Unique BGP Identifier**: NSX supports RFC6286, to relax the uniqueness of router ID between BGP peers across AS. BGP will form neighborship across AS even if the routers have the same router ID.

  Note This new behavior of accepting the same BGP RouterID will be the default in the system and cannot be changed.

How the BGP router ID (RID) is determined:

- If there is no loopback interface, BGP takes the highest interface IP address as RID.
- If BGP has already chosen the highest interface IP as RID, adding a loopback interface will not affect BGP neighborship and RID is not changed.
- If RID is the highest interface IP and loopback is present, disabling and enabling BGP will change the RID to the loopback IP.
- If RID is the highest interface IP and loopback is present, rebooting the edge node, enabling maintenance mode on the edge node, or restarting the routing process will not change the RID.
- If RID is the highest interface IP and loopback is present, redeploying or replacing the edge transport node will change the RID to the loopback interface IP.
- If RID is the highest interface IP and loopback is present, modifying or deleting the highest interface IP address will change the RID to the loopback interface IP.
- If RID is the loopback interface IP, modifying or deleting the highest interface IP will not change the RID.
- Clearing BGP neighbors will change the RID. It retains only the old RID.
- If the loopback interface has an IPv6 address, BGP does not use it as RID. It will take the highest IPv4 interface IP.
- A soft restart or hard restart of BGP adjacency from a remote site does not affect the BGP RID.

Supported BGP Capabilities

As defined in https://datatracker.ietf.org/doc/html/rfc2842, a BGP speaker determines the capabilities supported by its peer by examining the list of capabilities present in the Capabilities Optional Parameter in the OPEN message that the speaker receives from the peer. NSX supports the following capabilities:

Capability Code	Capability Description	Address Families Supported	Advertised Support from Tier-0 Gateway	Supported by Tier-0 Gateway when Received from Peer	Default Behavior	Configurable
1	Multiprotocol extensions, with: ■ AFI=1, SAFI=1 : IPv4 Unicast ■ AFI=2, SAFI=1 : IPv6 Unicast ■ AFI=25, SAFI=70 : L2VPN EVPN	IPv4 Unicast IPv6 Unicast L2VPN EVPN	Yes	Yes	IPv4 Unicast address family is enabled and advertised by default when an IPv4 neighbor is configured, or manually added in the Route Filter settings under the BGP neighbor configuration. IPv6 Unicast address family is enabled and advertised by default when an IPv6 neighbor is configured, or manually added in the Route Filter settings under the BGP neighbor configuration. L2VPN EVPN address family is enabled and advertised when configured in the Route Filter settings under the BGP neighbor configuration. The IPv4 Unicast address family is mandatory in NSX and automatically enabled when adding L2VPN EVPN address family.	Yes
2	Route refresh	IPv4 Unicast IPv6 Unicast L2VPN EVPN	Yes	Yes	Advertised by default	No
5	Extended next hop encoding	IPv6 Unicast	Yes	Yes	Not advertised by default. To enable this capability you must provide an IPv4 address family along with the IPv6 address family for the IPv6 BGP peer IP address.	Yes
64	Graceful restart	IPv4 Unicast IPv6 Unicast L2VPN EVPN	Yes	Yes	Not advertised by default (Edge node by default is a helper)	Yes

Capability Code	Capability Description	Address Families Supported	Advertised Support from Tier-0 Gateway	Supported by Tier-0 Gateway when Received from Peer	Default Behavior	Configurable
65	Support for 4-octet AS number	IPv4 Unicast IPv6 Unicast L2VPN EVPN	Yes	Yes	Advertised by default	No
69	ADD-Path, with: ■ AFI=1/2/25, SAFI=1/70 ■ Send/Receive=1 (Send Only) ■ Send/ Receive=2 (Receive Only) ■ Send/ Receive=3 (Both)	IPv4 Unicast IPv6 Unicast L2VPN EVPN	Yes (Receive only)	Yes (both Send and Receive)	The receive-only capability is supported and advertised by default. When the Edge node receives the same BGP prefix multiple times but with the same metric, if ECMP is enabled, all paths will be installed and active. When the Edge node receives the same BGP prefix multiple times with different metrics (for example, a larger ASPATH length) the best path route will be installed and active. The less preferred paths will be kept in the BGP routing table to improve control plane convergence.	No
73	FQDN	IPv4 Unicast IPv6 Unicast L2VPN EVPN	Yes	Yes	Advertised by default	No
128	Route refresh (Cisco)	IPv4 Unicast IPv6 Unicast L2VPN EVPN	Yes	Yes	Advertised by default	No

BGP Down Events

If a BGP down event occurs, an alarm will be raised. For details about the alarm, see the Routing Events table in NSX Event Catalog. In this release, additional information is provided regarding the reason for state changes of the BGP peer.

You can also use the PUT or PATCH method with the following APIs to disable a specific BGP peer. This will prevent alarms from being generated when that peer does not come up.

```
/policy/api/v1/infra/tier-0s/{tier-0-id}/locale-services/{locale-service-id}/bgp/neighbors/
{neighbor-id}
/policy/api/v1/global-infra/tier-0s/{tier-0-id}/locale-services/{locale-service-id}/bgp/
neighbors/{neighbor-id}
```

To disable a BGP peer, set the parameter `enabled` to `false`. For example:

```
PATCH https://<nsx-mgr>/policy/api/v1/infra/tier-0s/T0-1/locale-services/default/bgp/
neighbors/peer-1
{
  "enabled": "false"
}
```

You can also use the GET method to view the value of the `enabled` and other parameters.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Networking > Tier-0 Gateways**.

3 To edit a tier-0 gateway, click the menu icon (three dots) and select **Edit**.

4 Click **BGP**.

   a Enter the local AS number.

   In active-active mode, the default ASN value, 65000, is already filled in. In active-standby mode, there is no default ASN value.

   b Click the **BGP** toggle to enable or disable BGP.

   In active-active mode, **BGP** is enabled by default. In active-standby mode, **BGP** is disabled by default.

   c If this gateway is in active-active mode, click the **Inter SR iBGP** toggle to enable or disable inter-SR iBGP. It is enabled by default.

   If the gateway is in active-standby mode, this feature is not available.

   d Click the **ECMP** toggle button to enable or disable ECMP.

e Click the **Multipath Relax** toggle button to enable or disable load-sharing across multiple paths that differ only in AS-path attribute values but have the same AS-path length.

Note **ECMP** must be enabled for **Multipath Relax** to work.

f In the **Graceful Restart** field, select **Disable**, **Helper Only**, or **Graceful Restart and Helper**.

You can optionally change the **Graceful Restart Timer** and **Graceful Restart Stale Timer**.

By default, the Graceful Restart mode is set to **Helper Only**. Helper mode is useful for eliminating and/or reducing the disruption of traffic associated with routes learned from a neighbor capable of Graceful Restart. The neighbor must be able to preserve its forwarding table while it undergoes a restart.

For EVPN, only the **Helper Only** mode is supported.

The Graceful Restart capability is not recommended to be enabled on the tier-0 gateways because BGP peerings from all the gateways are always active. On a failover, the Graceful Restart capability will increase the time a remote neighbor takes to select an alternate tier-0 gateway. This will delay BFD-based convergence.

Note: Unless overridden by neighbor-specific configuration, the tier-0 configuration applies to all BGP neighbors.

5 Configure **Route Aggregation** by adding IP address prefixes.

a Click **Set** for Route Aggregation.

b Click **Add Prefix**.

c Enter a IP address prefix in CIDR format.

d For the option **Summary - Only**, select **Yes** or **No**.

6 Click **Apply**.

You must save the global BGP configuration before you can configure BGP neighbors.

7 Configure **BGP Neighbors**.

a Click **Set** for **BGP Neighbors**.

b Click **Add BGP Neighbor**.

c Enter the IP address of the neighbor.

d Enable or disable **BFD**.

e Enter a value for **Remote AS number**.

For iBGP, enter the same AS number as the one in step 4a. For eBGP, enter the AS number of the physical router.

f Under **Route Filter**, click **Set** to add one or more route filters.

For **IP Address Family**, you can select **IPv4**, **IPv6**, or **L2VPN EVPN**. The following combinations are supported:

- **IPv4** and **IPv6**
- **IPv4** and **L2VPN EVPN**

The combination of **IPv6** and **L2VPN EVPN** is not supported.

For the RFC 5549 feature, ensure that you provide an IPv4 address family along with the IPv6 address family for the IPv6 BGP peer IP address.

For **Out Filter** and **In Filter**, click **Configure** and select filters, then click **Save**.

For **Maximum Routes**, you can specify a value between **1** and **1,000,000**. When the number of BGP routes received from the peer reaches 75% of the configured limit, or when it exceeds the configured limit for the first time, a warning message is logged in the file `/var/log/frr.log`. In addition to the log message, the output of NSX CLI command `get bgp neighbor` shows if the number of prefixes received exceeds the configured limit. Note that the gateway will continue to accept routes from the BGP neighbor even after the **Maximum Routes** limit is reached.

Note If you configure a BGP neighbor with one address family, for example, **L2VPN EVPN**, and then later add a second address family, the established BGP connection will be reset.

g Enable or disable the **Allowas-in** feature.

This is disabled by default. With this feature enabled, BGP neighbors can receive routes with the same AS, for example, when you have two locations interconnected using the same service provider. This feature applies to all the address families and cannot be applied to specific address families.

h In the **Source Addresses** field, you can select a source address to establish a peering session with a neighbor using this specific source address. If you do not select any, the gateway will automatically choose one.

i Enter a value for **Max Hop Limit**.

j In the **Graceful Restart** field, you can optionally select **Disable**, **Helper Only**, or **Graceful Restart and Helper**.

Option	Description
None selected	The Graceful Restart for this neighbor will follow the Tier-0 gateway BGP configuration.
Disable	■ If the tier-0 gateway BGP is configured with **Disable**, Graceful Restart will be disabled for this neighbor. ■ If the tier-0 gateway BGP is configured with **Helper Only**, Graceful Restart will be disabled for this neighbor. ■ If the tier-0 gateway BGP is configured with **Graceful Restart and Helper**, Graceful Restart will be disabled for this neighbor.
Helper Only	■ If the tier-0 gateway BGP is configured with **Disable**, Graceful Restart will be configured as Helper Only for this neighbor. ■ If the tier-0 gateway BGP is configured with **Helper Only**, Graceful Restart will be configured as Helper Only for this neighbor. ■ If the tier-0 gateway BGP is configured with **Graceful Restart and Helper**, Graceful Restart will be configured as Helper Only for this neighbor.
Graceful Restart and Helper	■ If the tier-0 gateway BGP is configured with **Disable**, Graceful Restart will be configured as Graceful Restart and Helper for this neighbor. ■ If the tier-0 gateway BGP is configured with **Helper Only**, Graceful Restart will be configured as Graceful Restart and Helper for this neighbor. ■ If the tier-0 gateway BGP is configured with **Graceful Restart and Helper**, Graceful Restart will be configured as Graceful Restart and Helper for this neighbor.

Note For EVPN, only the **Helper Only** mode is supported.

k Enable **Configure Neighbor Local AS** if required.

- Choose 'Yes' if you want to override the Local AS number for this neighbor
- Choose 'No' if you do not want to override the Local AS number for this neighbor

Note **Neighbor Local AS** and **AS Path Modifier** fields appear after you enable **Configure Neighbor Local AS**.

l Enter the replaced local AS number, which you want to override the Local AS number for this neighbor, in **Neighbor Local AS**.

m Choose the AS path option in **AS Path Modifier**:

- **Default**: BGP prepends neighbor's local AS value to the AS path for both outgoing and incoming route advertisements from the peer neighbor. You can modify the default prepend action on the AS path in both inbound and outbound direction.
- **No Prepend**: The local router does not prepend the incoming advertisement from the peer with neighbor's local AS. The advertised AS path only prepends the BGP local AS of the router.
- **No Prepend Replace AS**: The local routes are advertised with the neighbor's local AS instead of the BGP's local AS to peer router.

n Click **Timers & Password**.

o Enter a value for **BFD Interval**.

The unit is milliseconds. For an Edge node running in a VM, the minimum value is 500. For a bare-metal Edge node, the minimum value is 50.

p Enter a value for **BFD Multiplier**.

q Enter a value, in seconds, for **Hold Down Time** and **Keep Alive Time**.

The **Keep Alive Time** specifies how frequently KEEPALIVE messages will be sent. The value can be between 0 and 65535. Zero means no KEEPALIVE messages will be sent.

The **Hold Down Time** specifies how long the gateway will wait for a KEEPALIVE message from a neighbor before considering the neighbor dead. The value can be 0 or between 3 and 65535. Zero means no KEEPALIVE messages are sent between the BGP neighbors and the neighbor will never be considered unreachable.

Hold Down Time must be at least three times the value of the **Keep Alive Time**.

r Enter a password.

This is required if you configure MD5 authentication between BGP peers.

8 Click **Save**.

9 (Optional) After a BGP neighbor is added, you can save its advertised and learned routes.

a Click the number from the **BGP Neighbors** field.

b From the **Set BGP Neighbors** dialog box, click the menu icon (3 dots) of a BGP neighbor and select **Download Advertised Routes** or **Download Learned Routes**.

BGP Autonomous System Number per Tier-0 VRF Gateway and BGP neighbor

Starting with NSX 4.1.0, you can configure a different BGP ASN per Tier-0 VRF gateway and per BGP neighbor.

BGP Autonomous System Number (ASN) per Tier-0 VRF Gateway

Tier-0 VRF gateways can have a different ASN other than the default/parent Tier-0 gateway local Autonomous System (AS). A separate ASN is important for Service Providers and multi-tenant topologies where end-customers provide their own BGP ASN.

When a Tier-0 VRF is created and configured for BGP, you can define an ASN within the VRF in the **Local AS** field. Else, you can leave the field blank to inherit the ASN from the parent Tier-0 gateway. To learn more, see Deploy VRF-Lite with BGP and Configure BGP.

Note BGP ASN per Tier-0 VRF gateway supports active/active and active/standby Edge Node HA modes. EVPN is currently not supported.

Changes in BGP ASN will cause the already established BGP connections to flap in the following situations until the session is re-established:

- When the parent Tier-0 gateway ASN changes, all child Tier-0 VRF gateways that inherit the ASN from the parent will change, causing all neighbors to flap.
- When the child Tier-0 VRF gateway has its own ASN different from the parent Tier-0 gateway:
  - If the parent Tier-0 gateway ASN is changed, the ASN of the child Tier-0 VRF gateway will not change, and no flap should occur.
  - If the child Tier-0 VRF gateway ASN is changed, all neighbors of that VRF will flap regardless of whether the neighbor has a local AS set. The parent Tier-0 gateway connections in other VRFs will not flap.
- Any change to the ASN for the neighbor local AS causes that connection to flap. No other neighbors should be affected.

BGP Per Neighbor ASN Override on Outgoing BGP Updates

NSX 4.1.0 provides BGP mechanisms for seamless ASN migration.

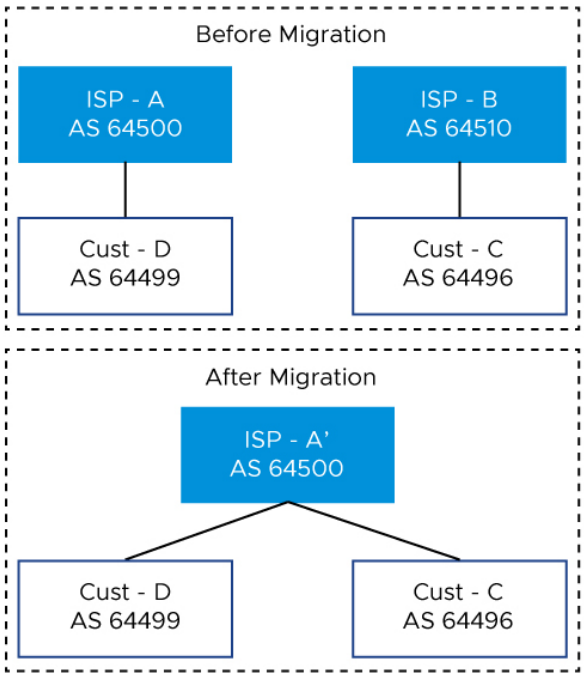
When an ISP acquires a network that belongs to a different autonomous system (AS), its BGP peers must be moved to the acquiring ISP's AS. This can be complicated and time consuming. Starting with NSX 4.1.0, with the BGP mechanisms, you can configure BGP neighbors in Tier-0 gateways and Tier-0 VRF gateways with a different local ASN than the gateway's ASN. Using a local ASN permits the routing devices in an acquired network to appear to belong to the former AS (Local AS).

BGP neighbors can be configured to have their ASN replaced in outgoing updates. Outgoing BGP updates use AS Path prepending. It can be used to influence the best path algorithm by making it longer or shorter, as well as to enable routers to migrate to new ASNs without having to upgrade all the peers simultaneously.

For more information, see RFC7795.

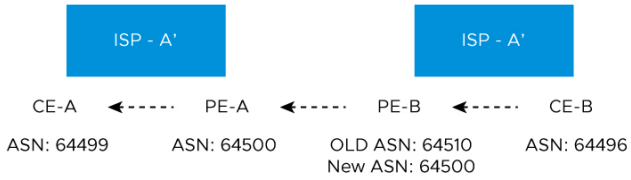
Note This is an EBGP feature. A neighbor's local-AS must be different from the gateway's local-AS. Changes in BGP local AS value will cause the already established neighbor BGP connections to flap until the session is re-established.

Let’s take an example based on merging two ISPs to one ASN of ISP A:



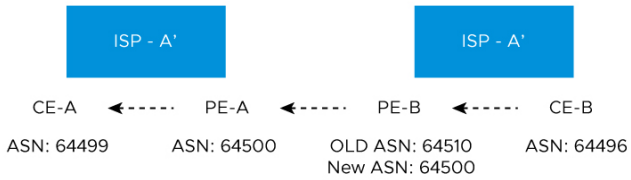
There three AS Path modifiers:

- **Default Prepend**: By default, AS Path prepends all received BGP updates with the old ASN while advertising it to other BGP speakers (A').



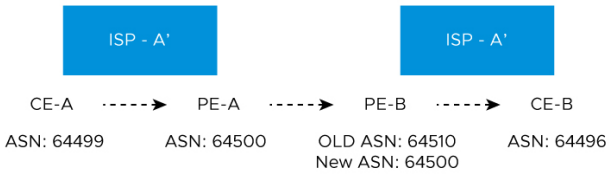
  A route received on CE-A from CE-B will have an AS PATH 64500 64510 64496.

- **No Prepend**: The local router will not prepend the old ASN to the incoming advertisement from that peer. AS Path will be prepended only with router's local AS.



  A route received on CE-A from CE-B will have an AS PATH 64500 64496.

- **No Prepend Replace AS**: Local routes will be advertised with the old ASN instead of the router's local AS, and they will appear to be running with the old ASN.



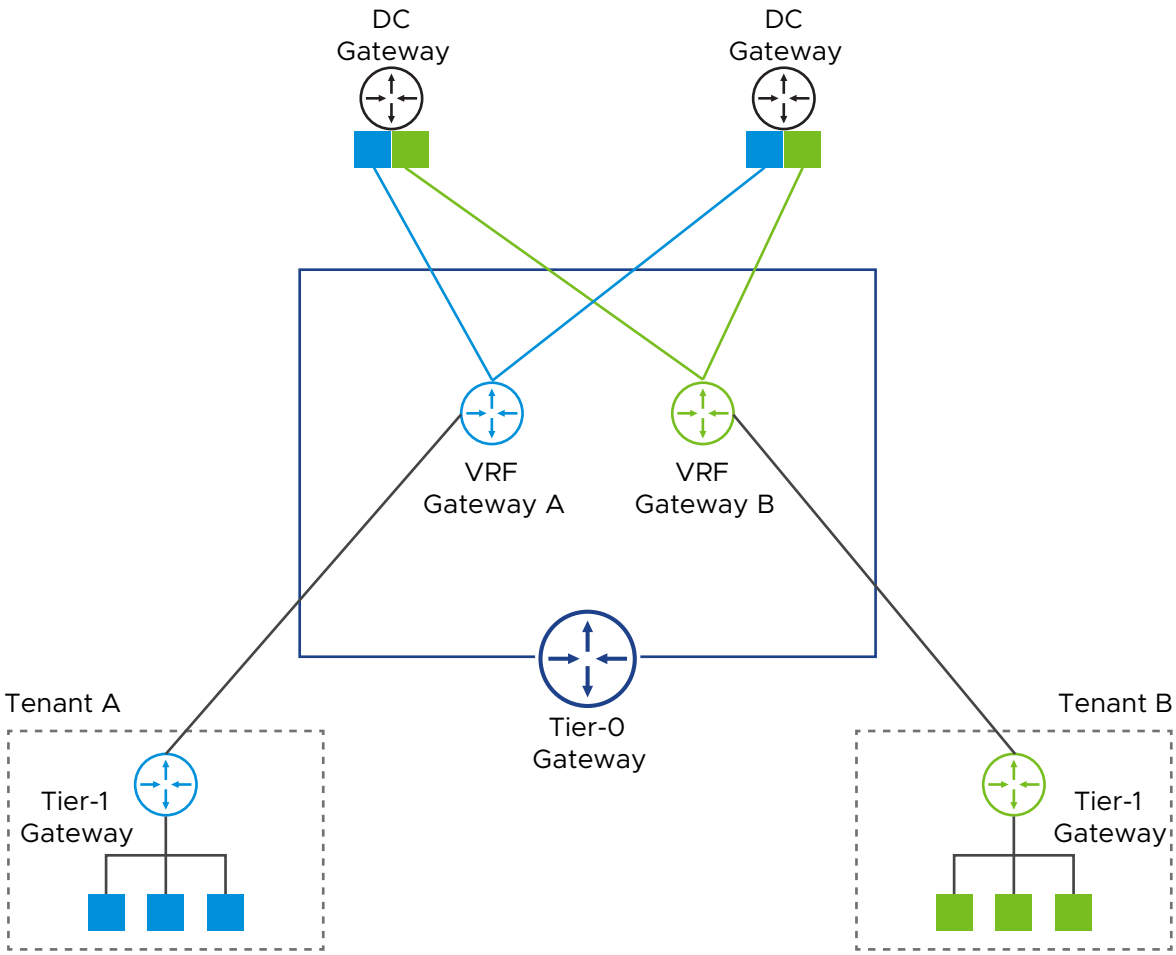
A route received on CE-B from CE-A will have an ASPATH 64510 64499.

Tier-0 VRF Gateways

Virtual routing and forwarding (VRF) makes it possible to instantiate isolated routing and forwarding tables within a router. VRFs are supported by deploying tier-0 VRF gateways. A tier-0 VRF gateway must be linked to a parent tier-0 gateway and inherits some of the tier-0 gateway settings, such as HA mode, Edge cluster, internal transit subnet, T0-T1 transit subnets, and BGP routing configuration.

Multiple tier-0 VRF instances can be created under the same parent tier-0, which allows the separation of segments and tier-1 gateways into multiple isolated tenants. With tier-0 VRF gateways, tenants can use overlapping IP addresses without any interference or communication with each other.

NSX tier-0 VRF gateways can be used to connect tenant networks to external routers using static routes or BGP [RFC4364]. This is also known as VRF-Lite.



NSX tier-0 VRF gateways can also be deployed with EVPN. For more information, see Chapter 12 Ethernet VPN (EVPN).

NSX Federation support:

- Tier-0 VRF gateway is not supported with NSX Federation and therefore it cannot be configured on Global Manager.
- Tier-0 VRF gateway is not supported on stretched tier-0 gateways in NSX Federation.

Note that even though a tier-0 VRF gateway has an HA mode, it does not have a mechanism to respond to a communication failure that is independent of the parent tier-0 gateway's mechanism. If a tier-0 VRF gateway loses connectivity to a neighbor but the criteria for the tier-0 gateway to fail over is not met, the VRF gateway will not fail over. The only time a VRF gateway will fail over is when the parent tier-0 gateway does a failover.

Deploy VRF-Lite with BGP

Prerequisites

- The parent tier-0 gateway needs to be created before the tier-0 VRF gateway instance.

- The parent tier-0 gateway needs to have an external interface before you create an external interface on the tier-0 VRF gateway.
- VLAN tagging (802.1q) is used to differentiate traffic among VRFs. The external interface on tier-0 VRF gateway needs to be connected to a trunk segment with the corresponding access VLAN ID defined in the segment VLAN range.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Configure the VLAN trunk segment.

   a Select **Networking > Segments**.

   b Click **Add Segments**.

   c Enter a name for the segment.

   d In **Connected Gateway**, set the type of connectivity for the segment as **None**.

   e Select a VLAN transport zone.

   f Expand the **Additional Settings** category.

   g In **VLAN**, enter a list or range of VLAN IDs allowed in the trunk segment.

   h Click **Save**.

3 Create the parent tier-0 gateway.

   The parent tier-0 gateway needs to be created before the tier-0 VRF gateway instance. For more information about configuring a tier-0 gateway, see Add a Tier-0 Gateway.

4 Create the tier-0 VRF gateway.

   a Select **Networking > Tier-0 Gateway**.

   b Click **Add Gateway > VRF**.

   c Enter a name for the gateway.

   d Select a tier-0 gateway in **Connect to Tier-0 Gateway**.

   Note Some advanced configurations are inherited from the parent tier-0, such as HA mode, edge cluster, internal transit subnet, T0-T1 transit subnets.

   e Click **VRF Settings**.

   Note The VRF settings are optional for regular VRF-Lite deployments, but are mandatory for EVPN use cases. For EVPN use cases, see Chapter 12 Ethernet VPN (EVPN).

f Under **L3 VRF Settings**, specify a **Route Distinguisher**.

If the connected tier-0 gateway has **RD Admin Address** configured, the **Route Distinguisher** is automatically populated. Enter a new value if you want to override the assigned Route Distinguisher.

g Click **Save** and then **Yes** to continue configuring the VRF gateway.

5 Configure the external interfaces on the VRF gateway.

a Click **Interfaces > Set > Add Interface**.

b Enter a name for the interface.

c Enter the IP address and mask for the external interface.

d In **Type**, select **External**.

e In **Connected To(Segment)**, select the trunk segment created from Step 2.

f Select an edge node.

g Enter the **Access VLAN ID** from the list as configured for the segment.

h Click **Save** and then **Close**.

6 Configure BGP neighbor for VRF-Lite.

a Click **BGP**.

b Click the **BGP** toggle to enable BGP.

c Enter **Local AS** number.

Note Leave this field blank to inherit the Local AS number from the parent tier-0 gateway.

You can configure the other advanced BGP settings such as ECMP.

d In the **BGP Neighbors** field, click **Set > Add BGP Neighbor**.

e Enter the neighbor IP address.

f Enable **BFD** if required.

g Enter the **Remote AS number** of the neighbor.

h Enter the source IP address.

There should be one or more addresses of created external interfaces or loopback.

i Under **Route Filter**, click **Set > Add Route Filter** to enable **IP Address Family**, filters based on prefix lists, and maximum routes received from the BGP neighbor.

j Enable or disable the **Allowas-in** feature.

This is disabled by default. With this feature enabled, BGP neighbors can receive routes with the same AS, for example, when you have two locations interconnected using the same service provider. This feature applies to all the address families and cannot be applied to specific address families.

k In the **Source Addresses** field, you can select a source address to establish a peering session with a neighbor using this specific source address. If you do not select any, the gateway will automatically choose one.

l Enter a value for **Max Hop Limit**.

m In the **Graceful Restart** field, you can optionally select **Disable**, **Helper Only**, or **Graceful Restart and Helper**.

Note For EVPN, only the **Helper Only** mode is supported.

n Enable **Configure Neighbor Local AS** if required.

- Choose 'Yes' if you want to override the Local AS number for this neighbor
- Choose 'No' if you do not want to override the Local AS number for this neighbor

Note **Neighbor Local AS** and **AS Path Modifier** fields appear after you enable **Configure Neighbor Local AS**.

o Enter the replaced local AS number, which you want to override the Local AS number for this neighbor, in **Neighbor Local AS**.

p Choose the AS path option in **AS Path Modifier**:

- **Default**: BGP prepends neighbor's local AS value to the AS path for both outgoing and incoming route advertisements from the peer neighbor. You can modify the default prepend action on the AS path in both inbound and outbound direction.
- **No Prepend**: The local router does not prepend the incoming advertisement from the peer with neighbor's local AS. The advertised AS path only prepends the BGP local AS of the router.
- **No Prepend Replace AS**: The local routes are advertised with the neighbor's local AS instead of the BGP's local AS to peer router.

q Click **Add** and then **Apply**.

r Click **Save** and then **Close**.

7 Re-distribute the routes in the VRF gateway and announce to the BGP neighbors.

a Click **Route Re-distribution**.

b In the **Route Re-distribution** field, click **Set > Add Route Re-distribution**.

c Enter a name for the redistribution policy.

d Click **Set** to select available sources, such as tier-0 connected interfaces and segments and then click **Apply**.

e Click **Add** and then click **Apply**.

8 Make sure that your segments or tier-1 gateways are connected to the tier-0 VRF gateway.

VRF Route Leaking

By default, the data plane traffic between VRF instances is isolated in NSX. By configuring VRF route leaking, traffic can be exchanged between VRF instances. Static routes must be configured on the tier-0 VRF instances to allow traffic to be exchanged.

Two topology options are supported in NSX:

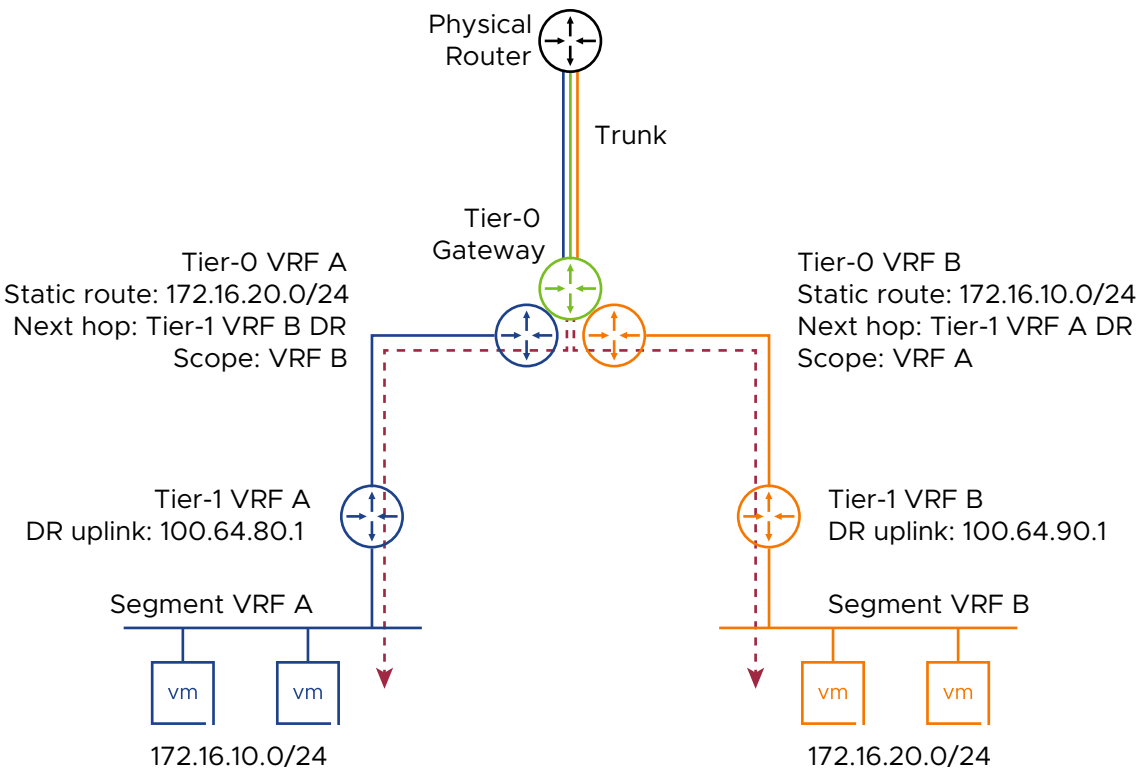
- Local VRF-to-VRF route leaking
- Northbound VRF leaking

Multi-tier routing architecture is required for traffic to be exchanged in a VRF leaking topology since static routes pointing to tier-1 distributed router (DR) uplinks are required.

Local VRF-to-VRF Route Leaking

Tier-1 DR IP addresses can be checked by using both the Edge node CLI or **Network Topology** in NSX Manager.

The following diagram depicts a sample topology for the local VRF-to-VRF route leaking option.



The configuration workflow for the sample topology is as follows:

Tier-0 VRF A Configuration Workflow

1. Select **Networking > Tier-0 Gateway**.
2. For tier-0 VRF A, click the menu icon (three dots) and select **Edit**.
3. Click **Routing**.
4. In the **Static Routes** field, click **Set > Add Static Route** and configure the static route:
   a. Enter a name for the static route.
   b. In the **Network** field, enter a prefix.

      For example, `172.16.20.0/24`
5. In the **Next Hops** column, click **Set > Set Next Hops** and define the next hops for the static route:
   a. Enter the IP address of the tier-1 DR uplink in VRF B.

      For example, `100.64.90.1`

      Note Tier-1 DR IP addresses can be checked by using both the Edge node CLI or **Network Topology** in NSX Manager.
   b. Enter the **Admin Distance** of `1`.
   c. Enter the scope.

      For example, `VRF-B`

Tier-0 VRF B Configuration Workflow

1. Select **Networking > Tier-0 Gateway**.
2. For tier-0 VRF B, click the menu icon (three dots) and select **Edit**.
3. Click **Routing**.
4. In the **Static Routes** field, click **Set > Add Static Route** and configure the static route:
   a. Enter a name for the static route.
   b. In the **Network** field, enter a prefix.

      For example, `172.16.10.0/24`
5. In the **Next Hops** column, click **Set > Set Next Hops** and define the next hops for the static route:
   a. Enter the IP address of the tier-1 DR uplink in VRF A.

      For example, `100.64.80.1`

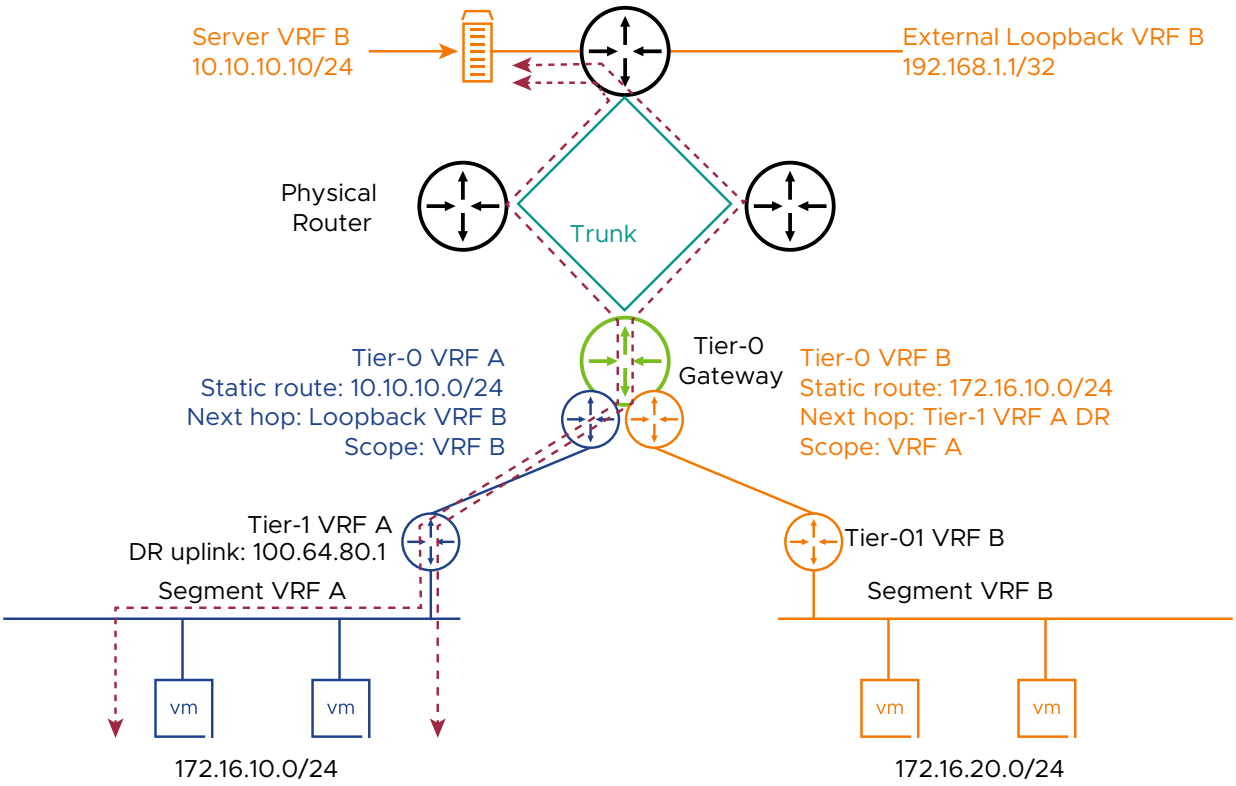
      Note Tier-1 DR IP addresses can be checked by using both the Edge node CLI or **Network Topology** in NSX Manager.
   b. Enter the **Admin Distance** of `1`.
   c. Enter the scope.

      For example, `VRF-A`

Northbound VRF Route Leaking

For this topology option, the northbound static route should have a next hop as an external IP address reachable in the destination VRF routing table. It is not recommended to point static routes directly to connected IP address uplinks as the static route would fail if an outage occurs on that link or neighbor. A loopback or virtual IP address host route (/32) can be advertised in the network in the destination VRF. Since the host route is advertised by both top of rack switches, two ECMP routes are installed in the tier-0 VRF. A return static route should be created in the destination VRF pointing to the tier-1 DR uplink IP address as the next hop.

The following diagram depicts a sample topology for the northbound VRF route leaking option.



The configuration workflow for the sample topology is as follows:

Tier-0 VRF A Configuration Workflow

1. Select **Networking > Tier-0 Gateway**.
2. For tier-0 VRF A, click the menu icon (three dots) and select **Edit**.
3. Click **Routing**.
4. In the **Static Routes** field, click **Set > Add Static Route** and configure the static route:
   a. Enter a name for the static route.
   b. In the **Network** field, enter a prefix.

      For example, `10.10.10.0/24`
5. In the **Next Hops** column, click **Set > Set Next Hops** and define the next hops for the static route:
   a. Enter the IP address of the tier-1 DR uplink in VRF B.

      For example, `192.168.1.1`
   b. Enter the **Admin Distance** of `1`.
   c. Enter the scope.

      For example, `VRF-B`

Tier-0 VRF B Configuration Workflow

1 Select **Networking > Tier-0 Gateway**.
2 For tier-0 VRF B, click the menu icon (three dots) and select **Edit**.
3 Click **Routing**.
4 In the **Static Routes** field, click **Set > Add Static Route** and configure the static route:
   a Enter a name for the static route.
   b In the **Network** field, enter a prefix.

   For example, `172.16.10.0/24`
5 In the **Next Hops** column, click **Set > Set Next Hops** and define the next hops for the static route:
   a Enter the IP address of the tier-1 DR uplink in VRF A.

   For example, `100.64.80.1`
   b Enter the **Admin Distance** of `1`.
   c Enter the scope.

   For example, `VRF-A`

Inter-VRF Routing

Starting with NSX 4.1.0, the Inter-VRF Routing feature allows you to import and export routes between VRFs.

You can import and export routes between VRFs using the following methods:

1 Inter-VRF import and export policies:

- Inter-VRF Route Advertisement - Advertise routes that are not BGP, such as Static, Connected, NAT, and all other sources that are available as inter-vrf static routes on connected gateway.
- Inter-VRF BGP Route Leaking - Enable BGP route leak for IPv4/IPv6 or both address families.

2 Manual Inter-VRF static route pointing to target VRF.

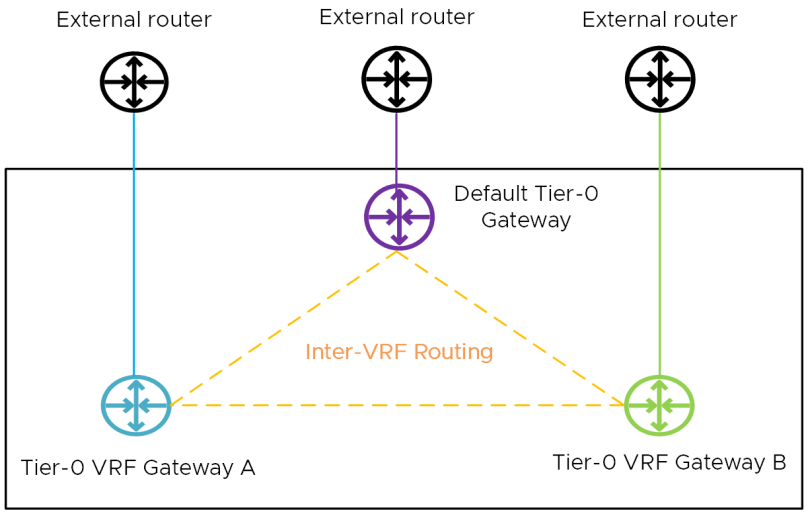
Inter VRF routing can be configured:

- From Default/Parent Tier-0 gateway and Tier-0 VRF gateway
- From Tier-0 VRF gateway to Default/Parent Tier-0 gateway
- From Tier-0 VRF gateway to Tier-0 VRF gateway

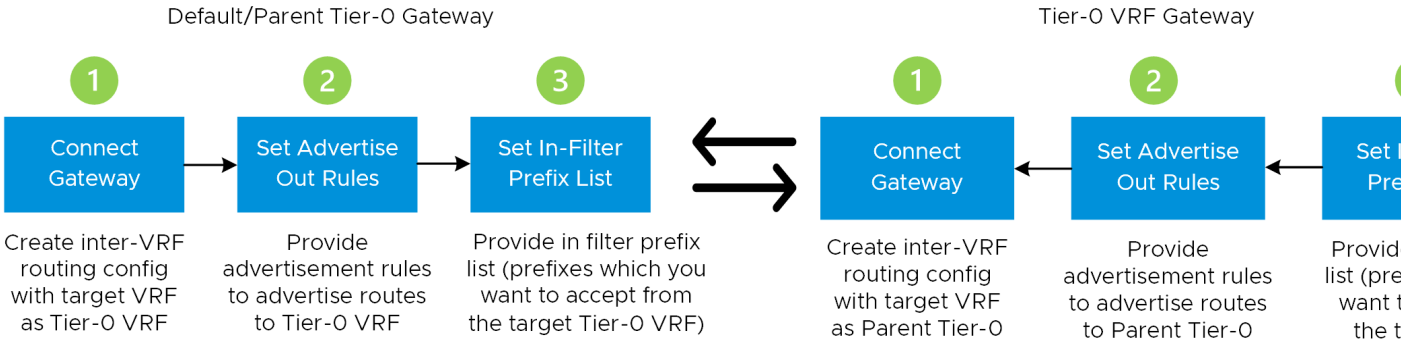
Configure Inter-VRF Route Advertisement

With Inter-VRF Route Advertisement, you can configure advertisement rules to advertise routes and prefixes on connected gateway to create inter-VRF routing between two VRFs.

Inter-VRF forwarding is enabled when Inter-VRF routing is configured on either end between two Tier-0 VRFs or between a Tier-0 VRF gateway and its Tier-0 gateway.



Configuration workflow:



You can enable Inter-VRF Route Advertisement by configuring:

- **Connect Gateway** - To attach the Target VRF to source.
- **Set Advertisement Out Rules** - To advertise the connected subnets of the source to the target VRF. You can allow or deny the subnets that you want to advertise to the target VRF i.e., you want the target VRF to learn.
- **Set In Filter Prefix Lists** - To select prefixes that you want to accept or deny from the target VRF.

Note NSX assigns an IPv4 and IPv6 address to each Tier-0 VRF gateway and Tier-0 gateway. IPv4 subnet is configured in the **Inter VRF Transit Subnet** field of the default Tier-0. For assigning IPv6 addresses to Inter-VRF interfaces, NSX allocates one IPv6 subnet. A /64 IPv6 subnet is allocated from the Tier-1-Tier-0 transit subnets configured on the default Tier-0.

Prerequisites

Ensure that the IP prefix list is created, see Create an IP Prefix List.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Connect Gateway:

a Select **Networking**.

b Select either the default/parent Tier-0 gateway or the Tier-0 VRF gateway, which you would like to configure.

c Click on the ellipsis button (three dots) and select **Edit**.

d Click **Routing**.

e Click **Set** next to **Inter VRF Routing**.

f Click **Add Inter VRF Routing**.

g In **Connected Gateway**, select the target VRF from the drop-down menu.

h Click **Save**.

Note Wait for the Status to change to 'Success'.

3 Set Advertisement Out Rules:

a Click on the ellipsis button (three dots) and select **Edit**.

b Click **Set** below **Advertisement Out Rules**.

c Click **Add Advertisement Out Rules**.

- Enter **Name** - Requires validation for Name.
- Enter **Subnets** - Specify subnets that are used to assign addresses to logical links connecting parent Tier-0 and child Tier-0 VRF.

  Note Leave the **Subnet** field blank to set it as 'any'. If subnet is 'any', Route Advertisement Type filters must be applied.

- Apply filter – Click on the (i) icon for more information about it. Toggle to 'Yes' if you would like to apply the filters.

  Note Disable the filters to advertise the configured aggregate prefix subnets.

- Advertisement Action - Action to advertise filtered routes to the connected Tier0 gateway.
  - Allow: Enables the advertisement
  - Deny: Disables the advertisement
- Prefix Operator - Prefix operator to filter subnets.
  - GE prefix operator: Filters all the routes with prefix length greater than or equal to the subnets configured.
  - EQ prefix operator: Filters all the routes with prefix length equal to the subnets configured.

  Note Prefix operation is ignored if no route advertisement type is specified.

- Select route advertisement types to be filtered.

d Click **Save**.

Note Wait for the Status to change to ‘Success.’

4 Set In Filter Prefix Lists:

a Click on the ellipsis button (three dots) and select **Edit**.

b Click **Set** below **In Filter Prefix Lists**

c Click **Add In Filter Prefix Lists**.

d Select the Prefix List that you had created in IP Prefix Lists window.

e Click **Add**.

f Click **Save**.

Note Wait for the Status to change to ‘Success’.

You may also download the advertised routes, by selecting **Download Advertised Network** from the ellipsis menu (three dots) next to the configured target VRF.

Configure Inter-VRF BGP Route Leaking

With Inter-VRF BGP Route Leaking, you can enable BGP route leak for IPv4/IPv6 or both address families.

Inter-VRF BGP route leaking uses BGP VPN to leak BGP routes. NSX internally configures RT(s) for route import and export.

Note Inter-VRF BGP Route Leaking is supported on both IPv4 and IPv6.

Prerequisites

- Enable BGP and provide local AS on both source and target VRF.
- On source VRF, configure redistribution rules or learn external routes through BGP neighbours (To add prefixes to BGP table). See Configure BGP.
- Use route-maps to import/export filter. See Create a Route Map.

Procedure

1. Select **Networking**.
2. Select either the default/parent Tier-0 Gateway or the Tier-0 VRF Gateway, which you would like to configure.
3. Click on the ellipsis button (three dots) and select **Edit**.
4. Click **Routing**.
5. Click **Set** next to **Inter VRF Routing**.
6. Click on the ellipsis button (three dots) and select **Edit**.
7. Click **Set** below **BGP Route Leaking**.
   a. Click on the ellipsis button (three dots) and select **Edit**.
   b. Toggle the **BGP Route Leaking** button to 'On' for IPv4 and IPv6.
   c. Assign a route map in the **In Filter** field for IPv4 and IPv6. The 'In filter' is used to determine which network or a range of networks are permitted or denied in the direction of In.
   d. Assign a route map in the **Out Filter** field for IPv4 and IPv6. The 'Out filter' is used to determine which network or a range of networks are permitted or denied in the direction of Out.
8. Click **Save**.

Results

Note BGP route leaking reduces the number of routes supported for the edge, since the BGP VPN table has to contain an additional copy of the route for import and export.

Create a Manual Inter-VRF Static Route with Next Hop as the Destination VRF.

You can create a manual inter-VRF static route with next hop as the destination VRF.

You can allow a datapath to route traffic to another VRF as the next hop scope using inter-vrf static routes. You can configure destination VRF path, for example 'Tier-0 VRF', as scope for route next hop. This will relax validation to require IP address, as internally it will translate the scope to the next hop IP.

Note The recommended way to configure inte-VRF route is to set scope as destination VRF and keep the next hop IP blank.

By default, static routes are tagged as 'INTER_VRF_STATIC', are not advertised to another VRF, and are not included in the BGP table. Static route only works if you have configured inter-VRF-routing with the destination VRF. Else, system will show an error.

Prerequisites

Inter-VRF-routing must be configured on both VRFs: source and destination.

Procedure

1. Select **Networking**.
2. Select either the default/parent Tier-0 Gateway or the Tier-0 VRF Gateway, which you would like to configure.
3. Click on the ellipsis button (three dots) and select **Edit**.
4. Click **Routing**.
5. Click **Set** next to **Static Routes**.
6. Click **Add Static Route** and configure the static route:
   a. Enter a name for the static route.
   b. In the **Network** field, enter a prefix.
7. Click **Set** below **Next Hops**.
8. Click **Set Next Hops** and define the next hops for the static route:
   a. Leave the IP address blank.
   b. Specify the administrative distance.
   c. Enter the destination VRF as scope.
9. Click **Add**.
10. Click **Apply**.

Note Wait for the Status to change to ‘Success’.

Configure OSPF

OSPF (Open Shortest Path First) is an interior gateway protocol (IGP) that operates within a single autonomous system (AS). Starting with NSX 3.1.1, you can configure OSPF on a tier-0 gateway.

The OSPF feature has the following capabilities and restrictions:

- Only OSPFv2 is supported.
- The tier-0 gateway can be active-active or active-standby (preemptive and non-preemptive).
- Only the default VRF is supported.
- You can configure a single area on a tier-0 gateway with a maximum of two tier-0 uplinks per Edge node.
- Backbone, normal area, and NSSA (not-so-stubby area) are supported.
- No redistribution is supported between BGP and OSPF.
- OSPF and BGP can be used together in the case of BGP multi-hop where the peer IP is learned through OSPF.
- The same redistribution features supported for BGP are supported for OSPF (tier-0 uplinks, downlinks, loopbacks, tier-1 downlinks, etc.). Depending on the area type, redistribution for all these networks will result in the Edge node generating type 5 external LSA (link-state advertisement) or type 7 external LSA with type 2 metric only (e2 or n2 routes). The Edge node itself can learn any type of LSA.
- MD5 and plain password authentication are supported on the area configuration.
- Federation is not supported.
- Route summarization for e2 and n2 routes is supported.
- The interface running OSPF can be broadcast or numbered point-to-point (/31).
- OSPF sessions can be backed with BFD.
- For graceful restart, only the helper mode is supported.
- Redistribution route maps are supported. Only the matching of prefix lists is applicable. No set actions.
- OSPF ECMP is supported up to maximum of 8 paths.
- Default Originate is supported.
- NAT with OSPF is not supported.
- Tier-0 VIP with OSPF is not supported.

How the OSPF router ID (RID) is determined:

- If there is no loopback interface, OSPF takes the highest interface IP address as RID.

- If OSPF has already chosen the highest interface IP as RID, adding a loopback interface will not affect OSPF neighborship and RID is not changed.
- If RID is the highest interface IP and loopback is present, disabling and enabling OSPF will change the RID to the loopback IP.
- If RID is the highest interface IP and loopback is present, rebooting the edge node, enabling maintenance mode on the edge node, or restarting the routing process will not change the RID.
- If RID is the highest interface IP and loopback is present, redeploying or replacing the edge transport node will change the RID to the loopback interface IP.
- If RID is the highest interface IP and loopback is present, modifying or deleting the highest interface IP address will change the RID to the loopback interface IP.
- If RID is the loopback interface IP, modifying or deleting the highest interface IP will not change the RID.
- Clearing OSPF neighbors will change the RID. It retains only the old RID.
- A soft restart or hard restart of OSPF adjacency from a remote site does not affect the OSPF RID.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Networking > Tier-0 Gateways**.

3 Click the **OSPF** toggle to enable OSPF.

4 In the **Area Definition** field, click **Set** to add an area definition.

You can add only one area definition.

a Click **Add Area Definition**.

b Enter an area ID.

The value must be a number or 4 numbers in IPv4 format (for example, 1.2.3.4).

c In the **Type** column, select **Normal** or **NSSA**.

An OSPF NSSA (not-so-stubby area) allows external routes to be flooded within the area.

d In the **Authentication** column, select **None**, **Password**, or **MD5**.

e In the **Key ID** column, enter a key ID if **Authentication** is set to **MD5**.

f In the **Password** column, enter a password if **Authentication** is set to **Password** or **MD5**.

In NSX 3.1.3.2 and earlier, the plain-text and MD5 passwords can have a maximum of 8 characters. Starting with NSX 3.1.3.3, the MD5 password can have a maximum of 16 characters, and the maximum length of the plain-text password remains to be 8 characters.

g Click **Save**.

5 In the **Graceful Restart** field, select either **Disable** or **Helper Only**.

6 Click the **ECMP** toggle to enable or disable ECMP.

ECMP (equal-cost multi-path) routing allows packet forwarding to occur over multiple best paths. It can provide fault tolerance for failed paths.

7 In the **Route Summarization** field, click **Set** to add a summary address.

Route summarization can reduce the number of LSAs that are flooded into an area. You can summarize one or more ranges of IP addresses and send routing information about these addresses in a single LSA.

a Click **Add Prefix**.

b Enter an IP address prefix in CIDR format.

c In the **Advertise** column, select **Yes** or **No** to indicate whether to advertise the summary route.

The default is **Yes**.

d Repeat the steps above to add more prefixes.

8 Click the **Default Route Originate** toggle to enable or disable default route originate.

Enable this to redistribute the default route in OSFP.

9 In the **OSPF Configured Interfaces** field, click **Set** to configure OSPF on existing external interfaces.

a Click **Configure Interface**.

b In the **Interface** column, select an interface from the dropdown list.

c In the **Area ID** column, select an area ID from the dropdown list.

d In the **Network Type** column, select **Broadcast** or **P2P**.

e In the **OSPF** column, set the toggle to **Enabled**.

f Click the **BFD** toggle to enable or disable BFD.

g If BFD is enabled, select a BFD profile.

h To change the **OSPF Hello Interval**, enter a new value.

The default is 10 seconds. This parameter specifies the time between Hello messages.

i To change the **OSPF Dead Interval**, enter a new value.

The default is 40 seconds. If a Hello message is not received within this time interval, the neighbor is considered unavailable..

j Click **Save**.

k Repeat the steps above to configure more interfaces.

10 Click **Save**.

11 Click **Route Re-distribution** to expand the section.

12 Click the **OSPF Route Redistribution Status** toggle to enable OSPF.

13 If you have route re-distribution rules configured, click the number to see the current rules or to add additional ones. If you do not have any configured, click **Set** to add re-distribution rules. Add **OSPF** to the **Destination Protocol** of any rule that will redistribute routes into OSPF. Remember to do this step if you plan to add re-distribution rules later.

Results

After you configure OSPF, in the **OSPF Neighbors** field, you can click **View** to see information about OSPF neighbors. The information displayed includes **Neighbor IP Address**, **Interface**, **Source**, **Edge Node**, **Priority**, and **State**.

Note: If a neighbor is not reachable, an alarm about the neighbor will be raised. If the neighbor is no longer in the network, simply acknowledge the alarm but do not resolve it. If you resolve the alarm, it will be raised again.

Configure BFD

BFD (Bidirectional Forwarding Detection) is a protocol that can detect forwarding path failures.

BFD can back up sessions for BGP and static routes.

Note In NSX 4.0.0.1, only IPv4 is supported. Starting with NSX 4.0.1.1, both IPv4 and IPv6 are supported.

Caution Note the following scenarios when there are connection failures involving BGP or BFD:

- With only BGP configured, if all BGP neighbors go down, the service router's state will be down.
- With only BFD configured, if all BFD neighbors go down, the service router's state will be down.
- With BGP and BFD configured, if all BGP and BFD neighbors go down, the service router's state will be down.
- With BGP and static routes configured, if all BGP neighbors go down, the service router's state will be down.
- With only static routes configured, the service router's state will always be up unless the node is experiencing a failure or in a maintenance mode.

For more information about BGP, see Configure BGP.

For more information about static routes, see Configure a Static Route.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Networking > Tier-0 Gateways**.

3 To edit a tier-0 gateway, click the menu icon (three dots) and select **Edit**.

4 Click **Routing** and **Set** for **Static Route BFD Peer**.

5 Click **Add Static Route BFD Peer**.

6 Select a **BFD profile**. See Add a BFD Profile.

7 Enter the peer IP address and optionally the source addresses.

8 Click **Save**.

Configure Multicast

IP multicast routing enables a host (source) to send a single copy of data to a single multicast address. Data is then distributed to a group of recipients using a special form of IP address called the IP multicast group address. You can configure multicast on a tier-0 gateway for an IPv4 network to enable multicast routing.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Networking > Tier-0 Gateways**.

3 To edit a tier-0 gateway, click the menu icon (three dots) and select **Edit**.

4 Click the **Multicast** toggle to enable multicast.

5 In the **Replication Multicast Range** field, enter an address range in CIDR format.

Replication Multicast Range is a range of multicast group addresses (GENEVE outer destination IP) that is used in the underlay to replicate workload/tenant multicast group addresses. It is recommended that there is no overlap between the Replication Multicast Range and workload/tenant multicast group addresses.

6 In the **IGMP Profile** drop-down list, select an IGMP profile.

7 In the **PIM Profile** drop-down list, select a PIM profile.

Configure IPv6 Layer 3 Forwarding

IPv4 layer 3 forwarding is enabled by default. You can also configure IPv6 layer 3 forwarding.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Networking > Networking Settings**.

3 Click the **Global Networking Config** tab.

4 Edit the Global Gateway Configuration and select **IPv4 and IPv6** for the **L3 Forwarding Mode**.

   IPv6 only is not supported.

5 Click **Save**.

6 Select **Networking > Tier-0 Gateways**.

7 Edit a tier-0 gateway by clicking the menu icon (three dots) and select **Edit**.

8 Go to **Additional Settings**.

   a There are no configurable IPv6 addresses for **Internal Transit Subnet**. The system automatically uses IPv6 link local addresses.

   b Enter a /48 IPv6 subnet for **T0-T1 Transit Subnets**.

9 Go to **Interfaces** and add an interface for IPv6.

Create SLAAC and DAD Profiles for IPv6 Address Assignment

When using IPv6 on a logical router interface, you can set up Stateless Address Autoconfiguration (SLAAC) for the assignment of IP addresses. SLAAC enables the addressing of a host, based on a network prefix advertised from a local network router, through router advertisements. Duplicate Address Detection (DAD) ensures the uniqueness of IP addresses.

Prerequisites

Navigate to **Networking > Networking Settings**, click the **Global Gateway Config** tab and select **`IPv4 and IPv6`** as the **L3 Forwarding Mode**

Procedure

1 From your browser, log in with admin privileges to an NSX Manager at https://<nsx-manager-ip-address>.

2 Select **Networking > Tier-0 Gateways**.

3 To edit a tier-0 gateway, click the menu icon (three dots) and select **Edit**.

4 Click **Additional Settings**.

5 To create an **ND Profile** (SLAAC profile), click the menu icon (three dots) and select **Create New**.

   a Enter a name for the profile.

   b Select a mode:

      - **Disabled** - Router advertisement messages are disabled.
      - **SLAAC with DNS Through RA** - The address and DNS information is generated with the router advertisement message.
      - **SLAAC with DNS Through DHCP** - The address is generated with the router advertisement message and the DNS information is generated by the DHCP server.
      - **DHCP with Address and DNS through DHCP** - The address and DNS information is generated by the DHCP server.
      - **SLAAC with Address and DNS through DHCP** - The address and DNS information is generated by the DHCP server. This option is only supported by NSX Edge and not by ESXi hosts.

   c Enter the reachable time and the retransmission interval for the router advertisement message.

   d Enter the domain name and specify a lifetime for the domain name. Enter these values only for the **`SLAAC with DNS Through RA`** mode.

   e Enter a DNS server and specify a lifetime for the DNS server. Enter these values only for the **`SLAAC with DNS Through RA`** mode.

   f Enter the values for router advertisement:

      - **RA Interval** - The interval of time between the transmission of consecutive router advertisement messages.
      - **Hop Limit** - The lifetime of the advertised routes.
      - **Router Lifetime** - The lifetime of the router.
      - **Prefix Lifetime**- The lifetime of the prefix in seconds.
      - **Prefix Preferred Time** - The time that a valid address is preferred.

6 To create a **DAD Profile**, click the menu icon (three dots) and select **Create New**.

   a Enter a name for the profile.

   b Select a mode:

      - **Loose** - A duplicate address notification is received but no action is taken when a duplicate address is detected.
      - **Strict** - A duplicate address notification is received and the duplicate address is no longer used.

c Enter the **Wait Time (seconds)** that specifies the interval of time between the NS packets.

d Enter the **NS Retries Count** that specifies the number of NS packets to detect duplicate addresses at intervals defined in **Wait Time (seconds)**

State Synchronization of Tier-0 Gateways

Connection information of the traffic running on a given tier-0 SR (Service Router) is synchronized to its peer tier-0 SR in active-standby or stateful active-active HA modes. Note that stateful active-active mode is only available starting with NSX 4.0.1.1.

In NSX 4.0.0.1, note the following about state synchronization:

- State synchronization is supported for Gateway Firewall, Identity Firewall, NAT, IPSec VPN, and DHCP.
- If new sessions were going through a tier-0 SR just before a failover, it might happen that those sessions were not synchronized on the associated tier-0 SR and potentially affect the traffic for those sessions.

Starting with NSX 4.0.1.1, note the following about state synchronization:

- In active-standby mode, state synchronization is supported for Gateway Firewall, Identity Firewall, NAT, IPSec VPN, and DHCP.
- In active-active mode, state synchronization is supported for Gateway Firewall, Identity Firewall, and NAT. IPSec VPN is not supported.
- If new sessions were going through a tier-0 SR just before a failover, it might happen that those sessions were not synchronized on the associated tier-0 SR and potentially affect the traffic for those sessions.

Changing the HA Mode of a Tier-0 Gateway

You can change the high availability (HA) mode of a tier-0 gateway in certain circumstances.

Changing the HA mode is allowed only if there is no more than one service router running on the gateway. This means that you must not have uplinks on more than one Edge transport node. However, you can have more than one uplink on the same Edge transport node.

After you set the HA mode from active-active to active-standby, you can set the failover mode. The default is non-preemptive.

HA mode change is not allowed if the following services or features are configured.

- DNS Forwarder
- IPSec VPN
- L2 VPN
- HA VIP
- Stateful Firewall

- SNAT, DNAT, NO_SNAT, or NO_DNAT
- Reflexive NAT applied on an interface
- Service Insertion
- VRF
- Centralized Service Port

Configure the ARP Limit of a Tier-0 or Tier-1 Gateway or Logical Router

You can configure the ARP limit of a tier-0 or tier-1 gateway or logical router using the API. The limit specifies the maximum number of ARP entries per transport node at each gateway or logical router.

To read or set the global ARP limit, use the following API methods and parameter:

Method	URI	Parameter
`GET`, `PUT`, `PATCH`	`/policy/api/v1/infra/global-config`	`arp_limit_per_gateway` (range: 5000 - 50000, default: 50000)

To read or set the ARP limit for a specific tier-0 or tier-1 gateway, use the following API methods and parameter. If the limit is not set, the global ARP limit will apply.

Method	URI	Parameter
`GET`, `PUT`, `PATCH`	`/policy/api/v1/infra/tier-0s/<tier-0-id>`	`arp_limit` (range: 5000 - 50000, no default)
`GET`, `PUT`, `PATCH`	`/policy/api/v1/infra/tier-1s/<tier-1-id>`	`arp_limit` (range: 5000 - 50000, no default)

Note that updating the ARP limit using Manager `GlobalConfig` API is not allowed.

Stateful Services on Tier-0 and Tier-1

To meet the demands of stateful services such as more bandwidth and throughput, you can configure Tier-0 and Tier-1 gateways in Active-Active (A-A) configuration. Stateful services are required for next generation firewall, Layer 7 rules, URL filtering or TLS decryption.

Starting with NSX 4.0.1.1, you can scale-out or scale-in the number of service routers by adding NSX Edge nodes to the cluster.

Caution As you scale-in or scale-out NSX Edge nodes, you might see loss of traffic packets for existing flows.

The supported stateful gateway services are:

- Gateway Firewall L3-L4

- APP-ID (L7)
- User-ID
- URL Filtering
- TLS Inspection
- IDS/IPS
- Malware Detection and Sandboxing
- NAT
- DHCP Relay Server
- DHCP Server

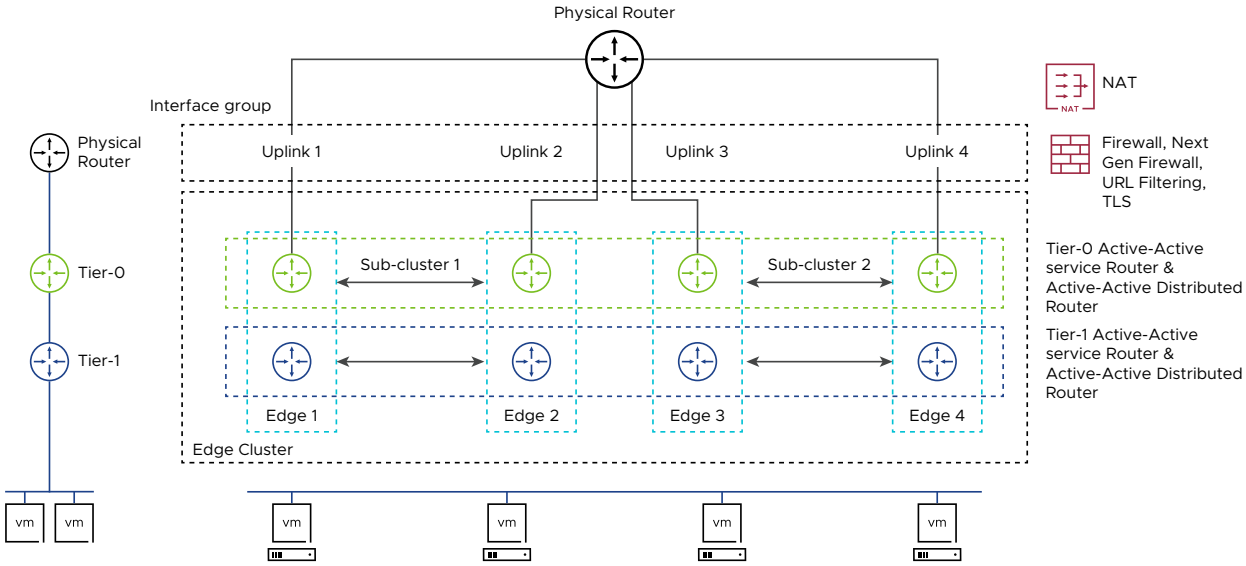
The unsupported services are:

- FQDN Analysis
- L2VPN
- IPSecVPN
- Gateway Network Introspection
- Local DHCP Server

In your existing topology, if Tier-1 gateway is in active-standby (A-S) mode, you cannot reconfigure it in A-A HA mode and it cannot share the same NSX Edge cluster with A-A stateful Tier-0 gateways. As a workaround, deploy that Tier-1 gateway in active-standby mode on a separate cluster. Then, deploy Tier-0 gateway on another NSX Edge cluster. If your environment requires a Tier-1 gateway, configure it in A-A HA mode. See Supported Topologies.

Key Concepts Stateful Services

Understand the key concepts required to configure stateful services.



Interface Groups

Use an interface group to group equivalent external (uplink) interfaces across service routers. Equivalent interfaces refer to interfaces that have the same inbound policies such as firewall rules, NAT rules, and so on, and equivalent network reachability. NSX only supports External interface. You must only create homogenous interface groups, such as an interface group comprising of only external interfaces.

To define an interface group, meet the following conditions:

- Only one interface link from each service router of the NSX Edge cluster must participate in the interface group.
- Every interface link must be part of a single interface group.
- Every interface group must have the same number of interfaces from each service router.

Note An interface group is required to create a stateful Active-Active group. If these conditions are not met, NSX raises a status alarm on the Tier-0 or the Tier-1 UI screen.

Interface groups allow traffic flows to continue without disruption even when the original uplink interface that supported packet transmission fails because the peer uplink interface takes over its traffic. So, if Uplink 1 or Edge 1 fails, then the interface group decides where the next packet must be punted to, which is the peer shadow port on the peer edge node on the same sub-cluster.

You can have more than one interface groups, where every interface group is dedicated to a specific requirement of the NSX Edge node. For example, one interface group can serve internet traffic while another group can be a Direct Connect connection between a NSX Edge cluster and a router.

On Tier-1 gateway, a default interface group is created. On Tier-0 gateway, you need to create an interface group.

To create an interface group, run the API `PUT /infra/tier-0s/<name>/locale-services/<location>/interface-groups/<group-name>`

```
{
    "resource_type": "Tier0InterfaceGroup",
    "id": "uplinkgroup",
    "display_name": "uplinkgroup",
    "path": "/infra/tier-0s/Tier0Gateway1/locale-services/Tier0LocalServices-1/interface-
groups/uplinkgroup",
    "relative_path": "uplinkgroup",
    "parent_path": "/infra/tier-0s/Tier0Gateway1/locale-services/Tier0LocalServices-1",
    "unique_id": "c5b2a758-7040-410b-a35d-298a16b55df0",
    "realization_id": "c5b2a758-7040-410b-a35d-298a16b55df0",
    "marked_for_delete": false,
    "overridden": false,
    "members": [
        {
            "interface_path": "/infra/tier-0s/Tier0Gateway1/locale-services/
Tier0LocalServices-1/interfaces/tier0_interface1"
        },
        {
            "interface_path": "/infra/tier-0s/Tier0Gateway1/locale-services/
Tier0LocalServices-1/interfaces/tier0_interface2"
        }
    ],
}
```

External Interface

Interface connecting to the physical infrastructure/router. It supports static routing and BGP. In previous releases, this interface was referred to as uplink interface. This interface can also be used to extend a VRF (Virtual routing and forwarding instance) from the physical networking fabric into the NSX domain.

Sub-clusters

When you configure stateful services on NSX Edge nodes, NSX automatically creates sub-clusters on the given NSX Edge cluster. So, NSX Edge a cluster of four NSX Edge nodes becomes two sub-clusters, where each sub-cluster is a pair of NSX Edge nodes.

All service routers on NSX Edge nodes participating in an interface group are converted into pairs.

For example, in sub-cluster 1, if Edge node 1 goes down, all ingress or egress traffic on Edge 1 is switched to Edge 2. So, Edge 1 and Edge 2 function as the original NSX Edge node and peer NSX Edge node respectively in the sub-cluster. During the failover process, Edge 2 takes over the backplane IP address that Edge 1 was serving to ensure no traffic is lost and traffic flow is maintained. When the failed Edge node 1 comes back up, the initial state is restored, where all traffic is redirected back to Edge 1.

Failure Domain

Configure failure domains to ensure that both NSX Edge nodes selected for a sub-cluster do not belong to the same failure domain.

To ensure failure domain functions as per design, meet these conditions:

- Label each NSX Edge with a failure domain and deploy one NSX Edge node in each failure domain. Both NSX Edge nodes of a sub-cluster must not belong to the same failure domain.
- Ensure that both NSX Edge nodes of a sub-cluster remain as part of the same sub-cluster. To ensure that these nodes are automatically paired in the same sub-cluster, follow a specific sequence when referencing these nodes to failure domains. For example, in a sub-cluster, first reference NSX Edge-1 to a failure domain and then reference NSX Edge-2 to a different failure domain. So, when NSX Edge-1 comes back up after a failure, the failure domain where the node was referenced to allows it to rejoin the same sub-cluster.

Supported Topologies

These are the supported topologies for stateful services on Tier-0 or Tier-1 in active-active HA mode.

Greenfield Topologies

In new installations, note the following considerations when building one of the supported topologies for stateful services on a NSX Edge cluster:

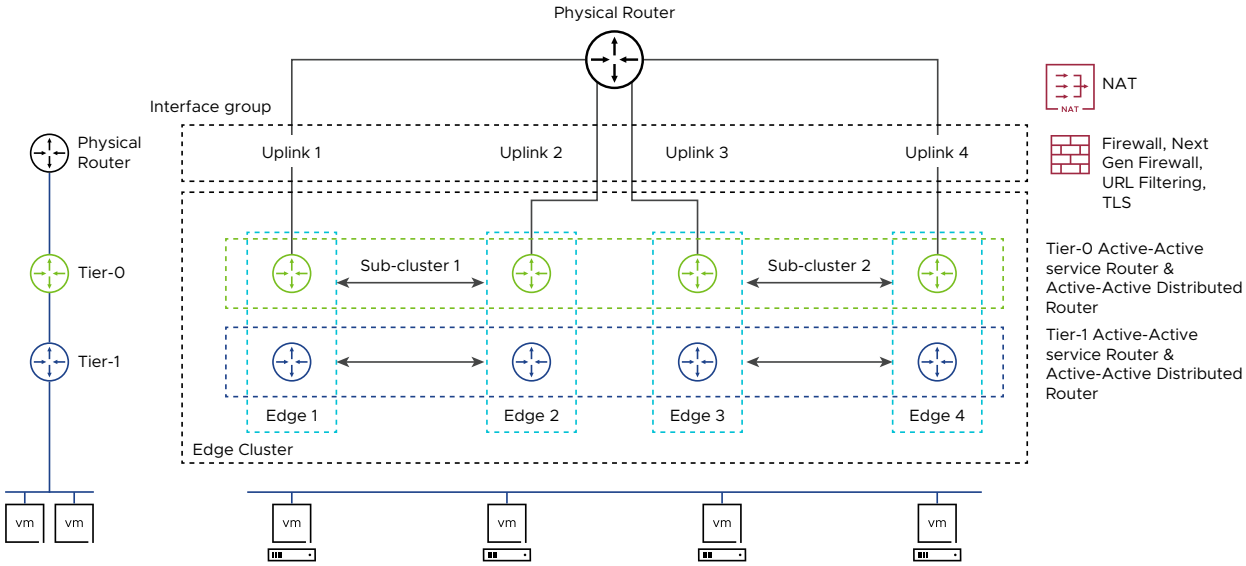
- Tier-1 sateful active-active gateways running stateful services must be connected to Tier-0 stateful active-active gateways and must be hosted on the same NSX Edge cluster.
- Tier-1 active-standby gateways can be connected to Tier-0 stateful active-active gateways but Tier-1 gateways must be hosted on a different NSX Edge cluster.

Brownfield Topologies

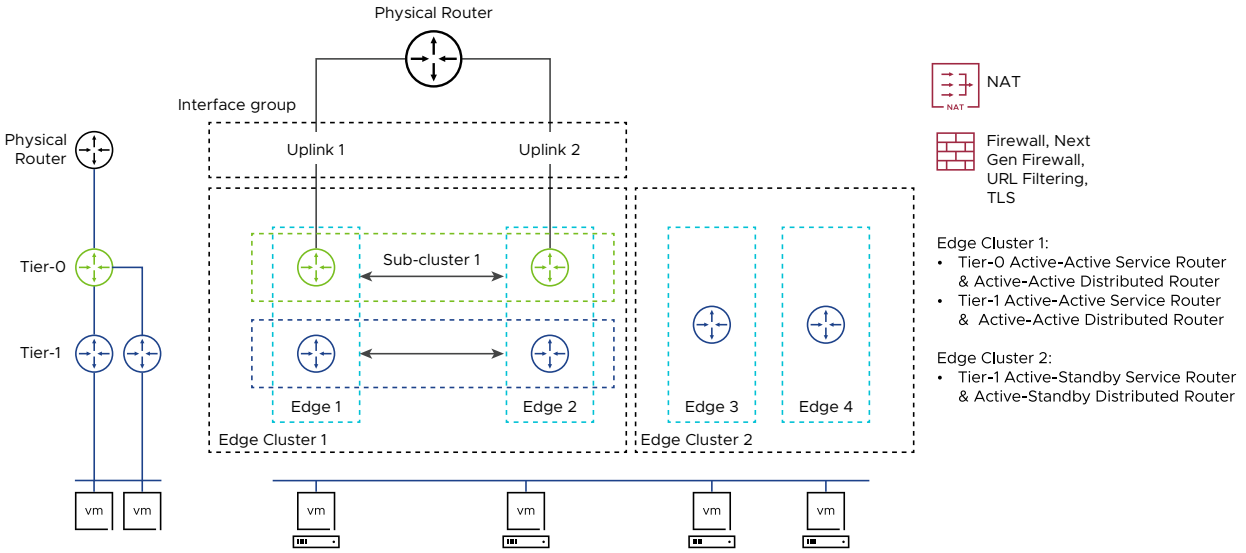
In existing installations, note the following considerations when building one of the supported topologies for stateful services on a NSX Edge cluster:

- An existing Tier-1 gateway in active-standby HA mode cannot be configured to be in active-active HA mode. You need to create a new Tier-1 gateway in active-active HA mode.
- Tier-0 active-standby gateways cannot be converted to Tier-0 active-active gateways.
- Tier-1 stateless active-active gateways can be converted to stateful active-active gateways if there no Tier-1 gateways attached to Tier-0 gateways.

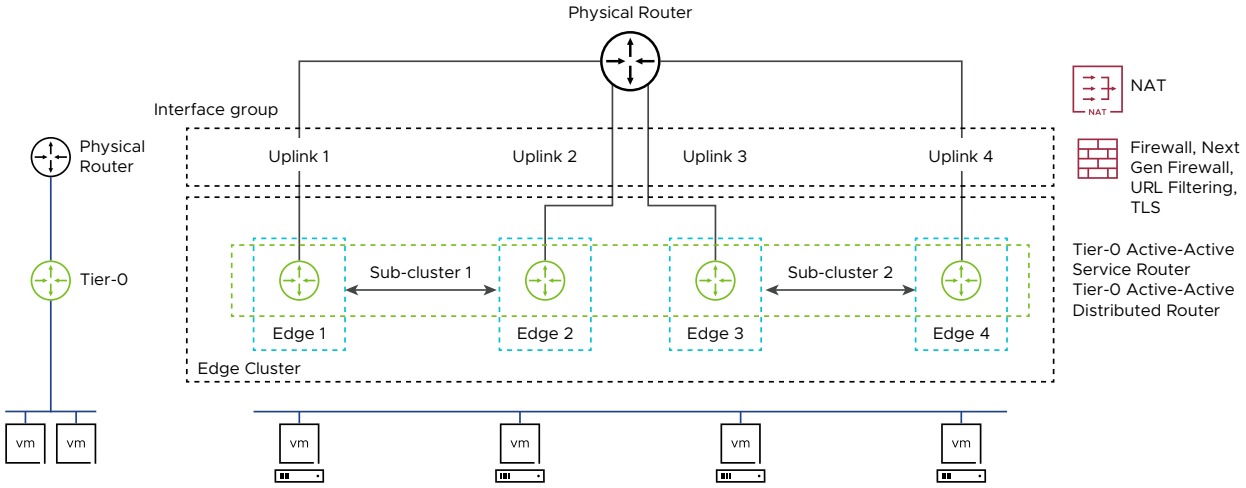
Tier-0 Active-Active and Tier-1 Active-Active HA mode



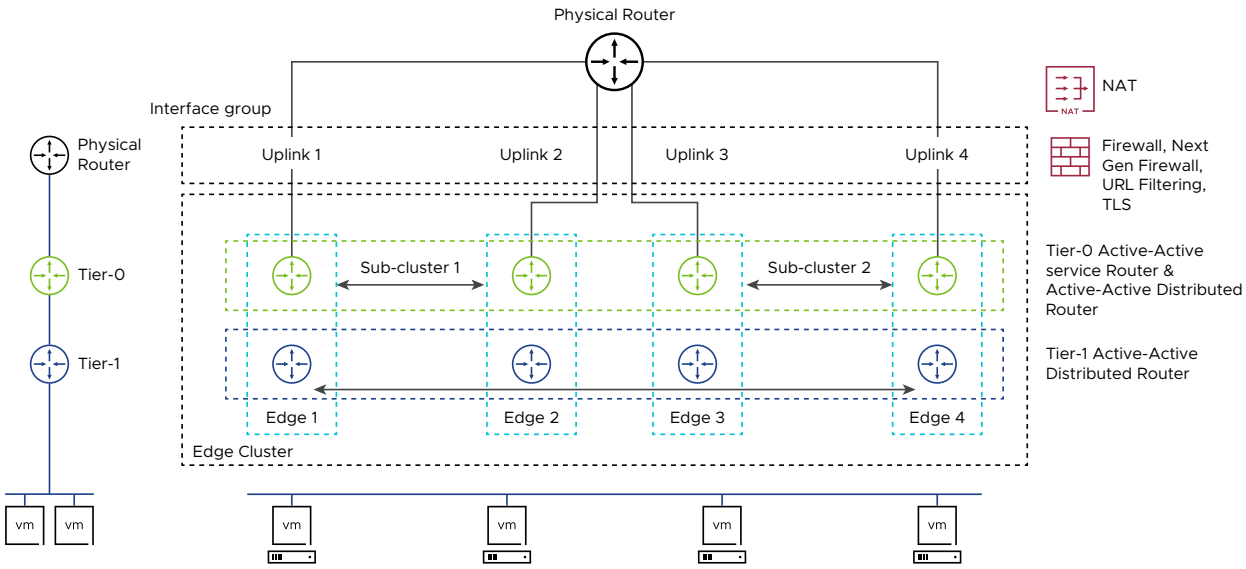
Tier-0 Active-Active and Tier-1 Active-Standby HA mode



Tier-0 Active-Active HA mode (no Tier-1 gateways)



Tier-0 Active-Active and Tier-1 Distributed Router only



Configure Failure Domains

A failure domain allows automatic recovery of a failed NSX Edge node, based on the allocation rules set in the NSX Edge cluster. Before configuring a Tier-0 stateful Active-Active (A-A) gateway, reference NSX Edge nodes to different failure domains.

A stateful A-A cluster expands or shrinks as you increase the number of NSX Edge nodes. In a stateful active-active cluster, NSX automatically creates sub-clusters out of the existing number of NSX Edge nodes. Each sub-cluster works as a pair of active and backup NSX Edge nodes. When one of the NSX Edge node from a sub-cluster fails, the failure domain associated to that NSX Edge node automatically recovers it.

In this procedure, you will reference NSX Edge nodes to different failure domains.

Note Ensure that NSX Edge-1 and NSX Edge-2 of sub-cluster-1 belong to two different failure domains.

Procedure

1 Using the API, create failure domains for the each Edge node that you will add to the stateful A-A cluster, for example, FD1A-Edge1 and FD2A-Edge 2. Set the parameter `preferred_active_edge_services` to true for both Edge 1 and Edge 2.

```
POST /api/v1/failure-domains
{
"display_name": "FD1A-Edge1",
"preferred_active_edge_services": "true"
}

POST /api/v1/failure-domains
{
"display_name": "FD2A-Edge2",
"preferred_active_edge_services": "true"
}
```

2 Using the API, associate each Edge node with the failure domain for the site. First call the `GET /api/v1/transport-nodes/<transport-node-id>` API to get the data about the Edge node. Use the result of the GET API as the input for the `PUT /api/v1/transport-nodes/<transport-node-id>` API, with the additional property, `failure_domain_id`, set appropriately. For example,

```
GET /api/v1/transport-nodes/<transport-node-id>
Response:
{
    "resource_type": "TransportNode",
    "description": "Updated NSX configured Test Transport Node",
    "id": "77816de2-39c3-436c-b891-54d31f580961",
    ...
}
```

```
PUT /api/v1/transport-nodes/<transport-node-id>
{
    "resource_type": "TransportNode",
    "description": "Updated NSX configured Test Transport Node",
    "id": "77816de2-39c3-436c-b891-54d31f580961",
    ...
    "failure_domain_id": "<UUID>",
}
```

3 Using the API, configure the Edge cluster to allocate nodes based on failure domain. First call the `GET /api/v1/edge-clusters/<edge-cluster-id>` API to get the data about the Edge cluster. Use the result of the GET API as the input for the `PUT /api/v1/edge-clusters/<edge-cluster-id>` API, with the additional property, `allocation_rules` set appropriately. For example,

```
GET /api/v1/edge-clusters/<edge-cluster-id>
Response:
{
    "_revision": 0,
    "id": "bf8d4daf-93f6-4c23-af38-63f6d372e14e",
    "resource_type": "EdgeCluster",
    ...
}
```

```
PUT /api/v1/edge-clusters/<edge-cluster-id>
{
    "_revision": 0,
    "id": "bf8d4daf-93f6-4c23-af38-63f6d372e14e",
    "resource_type": "EdgeCluster",
    ...
    "allocation_rules": [
        {
            "action": {
                      "enabled": true,
                      "action_type": "AllocationBasedOnFailureDomain"
                      }
        }
    ],
}
```

Results

The NSX Edge nodes are referenced to different failure domains. You can now use them to create a cluster and configure Tier-0 gateway in A-A Stateful HA mode.

Configure Stateful Services on Tier-0 and Tier-1 Gateways

Configure Tier-0 and Tier-1 gateways in Active-Active (A-A) Stateful high availability mode on an NSX Edge cluster and enable stateful services.

The topology considered for this procedure uses Tier-0 gateways and Tier-1 gateways both in A-A Stateful mode.

Prerequisites

- If there are odd number of NSX Edge nodes in the cluster, it leads to the scenario where one sub-cluster does not have a backup node. On failure of that single node, traffic is disrupted.

NSX triggers an alarm that you must resolve to correctly configure stateful services. Ensure a NSX Edge cluster consists of even number of NSX Edge nodes. For example, in a NSX Edge cluster of 4 nodes, NSX forms two sub-clusters, where each sub-cluster contains two nodes. One node in each sub-clusters is the backup node of the active NSX Edge node.

- Ensure the NSX Edge nodes you will use as part of the NSX Edge cluster are referenced to different failure domains.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Go to **Networking** → **Tier-0 Gateways**.

3 From the **Add Gateway** drop-down menu, click **Tier-0**.

4 Enter the name of the Tier-0 gateway.

5 In the HA Mode field, select **Active Active** and enable **Stateful**.

Note Once you enable the gateway to be stateful, you cannot edit the HA mode.

6 Select the NSX Edge cluster and click **Save**.

7 Click **Yes** to continue to edit the Tier-0 gateway.

8 Expand the **Interface and Interface Groups** section and in the **External** field click **Set**.

9 In the **Set Interfaces** window, click **Add Interface**.

10 Enter the name, select the segment the interface is connected to and NSX Edge node. Enter any other optional details.

11 Click **Save** to complete adding the interface.

12 After you add interfaces, go to the **Interface Groups** field, click **Set**.

13 In the **Set Interface Groups** window, click **Add Interface Group**.

Important Create an interface group comprising of one uplink from each Tier-0 SR of the cluster. Ensure that one uplink from every SR is part of the group and that uplink is only part of a single group. Each interface participating in the interface group must be equivalent. Uplinks are called equivalent when they are reachable on the network and when they share the same firewall, NAT and other network layer 4-7 policies.

An interface group allows multiple segments to be grouped into a single group which is connected to a NSX Edge cluster.

If the interface group does not have an uplink from each SR, then it can result in traffic loss. NSX triggers an alarm when this requirement is not met.

14 Click **Close Editing** to update the Tier-0 A-A HA gateway.

15 After deploying the Tier-0 A-A Stateful gateway, deploy Tier-1 gateways in A-A HA mode on the same NSX Edge cluster where Tier-0 gateways are configured. When you scale-out or scale-in a Tier-0 gateway, which means new sub-clusters of NSX Edge are added or removed, associated Tier-1 gateways also follow the same behavior.

16 Create a locale service on Tier-0 gateways.

`PUT https://<policy-mgr>/policy/api/v1/infra/tier-0s/vmc_prv/locale-services/<locale_service>`

```
{
  "route_redistribution_types": [ "TIER0_STATIC", "TIER0_NAT" ],
  "edge_cluster_path": "/infra/sites/default/enforcement-point/nsx/edge-clusters/
<95196903-6b8a-4276-a7c4-387263e834fd>",
  "preferred_edge_paths": [ "/infra/sites/default/enforcement-point/nsx/edge-clusters/
<95196903-6b8a-4276-a7c4-387263e834fd>/edge-nodes/<940f1f4b-0317-45d4-84e2-b8c2394e7405>"],
  "_revision": 0
}
```

17 Deploy Tier-1 gateways in A-A HA mode and select the NSX Edge cluster to run the gateway.

18 Create a locale service on Tier-1.

Without creating a locale service, the gateway is a DR-only gateway.

`PUT https://<policy-mgr>/policy/api/v1/infra/tier-1s/cgw/locale-services/<locale_service>`

```
{
  "edge_cluster_path": "/infra/sites/default/enforcement-point/nsx/edge-clusters/
<95196903-6b8a-4276-a7c4-387263e834fd>",
  "preferred_edge_paths": [ "/infra/sites/default/enforcement-point/nsx/edge-clusters/
<95196903-6b8a-4276-a7c4-387263e834fd>/edge-nodes/<940f1f4b-0317-45d4-84e2-b8c2394e7405>"],
  "_revision": 0
}
```

19 Create an SNAT rule for the service router on the Tier-0 A-A Stateful gateway. It is mandatory to enter the translated IP.

20 Go to **Networking** → **NAT** and click **Add NAT Rule**.

21 From the **Action** drop-down list, select **SNAT** and enter source and destination IP.

22 In the **Translated IP | Port** field, enter the IP that the source IP must be translated to.

23 Click **Save**.

24 Verify the high availability status on Tier-1 SR and Tier-0 SR. Verify that a pair of NSX Edge nodes form a sub-cluster. Both are active. The peer node only takes over and processes traffic on the failure of the active node.

On a Tier-0 node> `# get high-availability status`

```
Service Router
UUID            : 073a9fda-7a11-4d59-80c3-a7ea5371d265
```

```
state           : Active
type            : TIER0
mode            : Stateful A/A
failover mode   : Preemptive
rank            : 0
service count   : 0
service score   : 0
HA ports state
    UUID        : de647a80-d27c-46ee-a251-b35a3cead0d0
    op_state    : Up
    addresses   : 169.254.0.2/25;fe80::50:56ff:fe56:5300/64
Sub-cluster Information
    UUID            : c8db92e7-21da-453d-9853-2648849e7bda
    Peer SR UUID    : daaca25b-9028-4e31-b9b7-35bae481e60a
    Peer Node UUID  : 68668f1c-0330-11ec-84cf-00505682699c
Peer Routers
    Node UUID   : 9fe732b6-0330-11ec-ae4e-005056821b5a
    HA state    : Active
    Node UUID   : 8486560a-0330-11ec-902b-00505682411d
    HA state    : Active
    Node UUID   : 68668f1c-0330-11ec-84cf-00505682699c
    HA state    : Active
```

Results

You can run stateful services on Tier-0 gateways in active-active mode.

What to do next

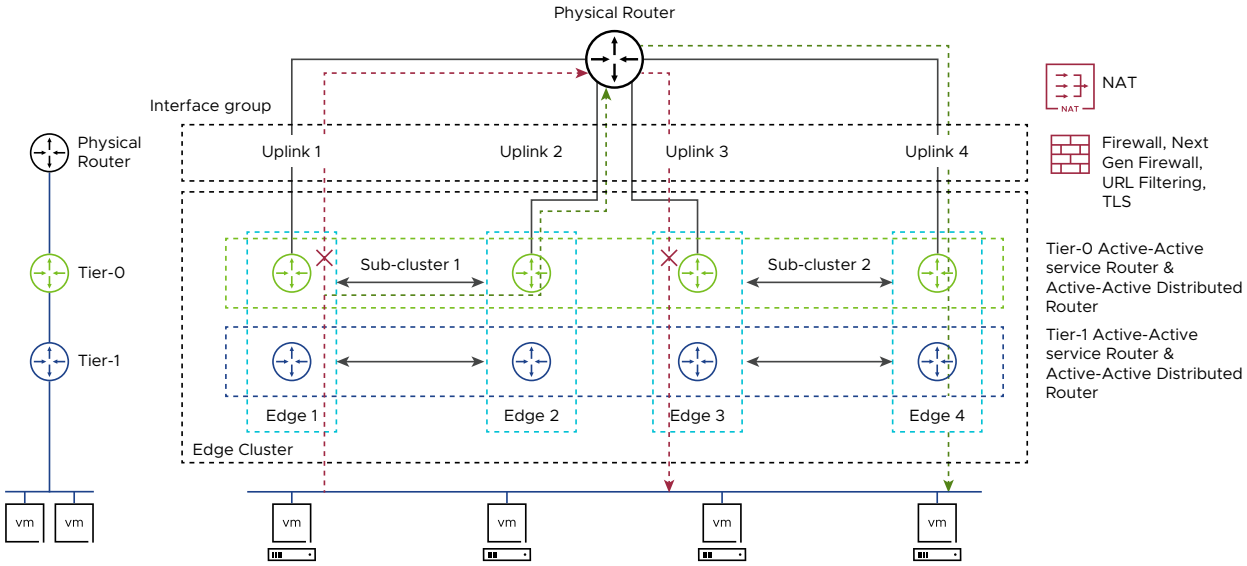
- To scale-out a cluster, add even number of NSX Edge nodes.

  Note If you add odd number of NSX Edge nodes, the newly added node does not have a backup node. If the newly added node fails, then traffic is disrupted. The NSX Manager raises an alarm if you only add odd number of nodes in a NSX Edge cluster.

- To scale-in a cluster, remove even number of NSX Edge nodes from the cluster.

Understanding Traffic Flows

Traffic flows on stateful Tier-0 and Tier-1 gateways configured in active-active HA mode.



South-North Traffic Flow

1. Based on a deterministic hash, an incoming packet from a southbound VM is punted to the backplane of the Edge-2.
2. Edge-2 determines that Edge-4 is actively managing the traffic flows and forwards the flow out through the external interfaces (which are part of the interface group).
3. An IP hash is performed, based on external server destination IP, and traffic is punted from Edge-2 to Edge-4. The packet is further forwarded to Tier-0 gateway service router (SR), where SNAT changes the source IP address to translated IP address.
4. After the flow reaches Edge-4 Tier-0 SR, the shadow port forwards the NAT traffic to the uplink interface and then sent out to the physical router.
5. If Tier-0 SR on Edge-4 fails, NSX punts traffic to its backup node in the sub-cluster, Edge-3, where SNAT changes the source IP address to translated IP address. The backup interface on Edge-1 takes over the backplane IP and the uplink IP of Tier-0 gateway before beginning to process traffic. The backup interface on Edge-3 is operationally `Up` and the shadow interface on Edge-4 is `Down`.
6. All traffic flows processed by firewall and NAT rules are synchronized on the Tier-0 SR on Edge-3.
7. When Edge-4 comes back up, the flow is resynchronized back to Edge-4. When the shadow port comes back up, NSX punts traffic to it.

North-South Traffic Flow

1. A packet from a northbound VM is hashed by the physical router using its own hashing algorithm to send the packet to Edge-3, based on an ECMP routing choice. The Tier-0 gateway is running on Edge-3.

2 Edge-3 determines that Edge-4 is actively managing the traffic flow and forwards the flow to Edge-4. The flow is managed by the shadow interface of Edge-4.

3 An IP hash is performed, based on external server source IP, traffic is punted from Edge-3 Tier-0 SR to Edge-4 Tier-0 SR, where NAT is enabled. The source IP is changed to the translated IP address.

4 The packet is sent from Edge-4 Tier-0 SR to Edge 4 Tier-0 DR and then to Tier-1 gateway, finally reaching the destination VM.

5 If Tier-0 service router on Edge-4 fails, NSX punts traffic to its peer node (sub-cluster 2), which is Edge-3. NAT enabled on Edge-3 changes the source IP address to translated IP address.

6 Before beginning to process traffic, the backup shadow port on Edge-3 manages the traffic flow. Now, the backup shadow port on Edge-3 is operationally `Up` and the shadow port on Edge-4 is `Down`.

7 All traffic flows processed by firewall and NAT rules are synchronized on the Tier-0 SR on Edge-3.

8 When Edge-4 comes back up, the flow is resynchronized back to it. The shadow port on Edge-4 comes back up and manages the punted traffic.

Sub-cluster Failure

If both the nodes in a sub-cluster go down, the sub-cluster goes down.

- Existing flows are disrupted causing traffic loss.
- New flows are punted to the other sub-cluster.
- When the failed sub-cluster comes back up again, the flows return to the original sub-cluster.

If a sub-cluster goes down for any reason, then the other sub-cluster in the cluster takes over.

Single Node Failure

On failure of an Edge node , the following events happen:

1 Interface links of the Edge node fail.

2 The shadow port on the failed Edge node is in `Down` state.

3 The backup port of the peer node in the sub-cluster takes over.

4 The firewall and the NAT states are synchronized on the peer Edge node.

5 The backup port on the peer node provides connectivity to new traffic flows.

6 When interface links of the failed node comes back up, the firewall and the NAT states are resynchronized with the shadow port of the active node.

7 NSX punts back traffic flows to the original node.

Tier-1 Gateway

3

A tier-1 gateway has downlink connections to segments and uplink connections to tier-0 gateways.

You can configure route advertisements and static routes on a tier-1 gateway. Recursive static routes are supported.

This chapter includes the following topics:

- Add a Tier-1 Gateway
- State Synchronization of Tier-1 Gateways

Add a Tier-1 Gateway

A tier-1 gateway is typically connected to a tier-0 gateway in the northbound direction and to segments in the southbound direction.

If you are adding a tier-1 gateway from Global Manager in NSX Federation, see Add a Tier-1 Gateway from Global Manager.

Tier-0 and tier-1 gateways support the following addressing configurations for all interfaces (external interfaces, service interfaces and downlinks) in both single tier and multi-tiered topologies:

- IPv4 only
- IPv6 only
- Dual Stack - both IPv4 and IPv6

To use IPv6 or dual stack addressing, enable **IPv4 and IPv6** as the L3 Forwarding Mode in **Networking > Networking Settings > Global Networking Config**.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Select **Networking > Tier-1 Gateways**.
3. Click **Add Tier-1 Gateway**.
4. Enter a name for the gateway.

5 (Optional) Select a tier-0 gateway to connect to this tier-1 gateway to create a multi-tier topology.

6 (Optional) Select an NSX Edge cluster if you want this tier-1 gateway to host stateful services such as NAT, load balancer, or firewall.

If an NSX Edge cluster is selected, a service router will always be created (even if you do not configure stateful services), affecting the north/south traffic pattern.

7 (Optional) In the Edges field, click **Set** to select an NSX Edge node.

8 If you selected an NSX Edge cluster, select a failover mode or accept the default.

Option	Description
Preemptive	If the preferred NSX Edge node fails and recovers, it will preempt its peer and become the active node. The peer will change its state to standby.
Non-preemptive	If the preferred NSX Edge node fails and recovers, it will check if its peer is the active node. If so, the preferred node will not preempt its peer and will be the standby node. This is the default option.

9 If you plan to configure a load balancer on this gateway, select an **Edges Pool Allocation Size** setting according to the size of the load balancer.

The options are **Routing**, **LB Small**, **LB Medium**, **LB Large**, and **LB XLarge**. The default is **Routing** and is suitable if no load balancer will be configured on this gateway. This parameter allows the NSX Manager to place the tier-1 gateway on the Edge nodes in a more intelligent way. With this setting the number of load balancing and routing functions on each node is taken into consideration. Note that after you create the gateway, you can change this setting if you have not configured a load balancer.

10 (Optional) Click the **Enable StandBy Relocation** toggle to enable or disable standby relocation.

Standby relocation means that if the Edge node where the active or standby logical router is running fails, a new standby logical router is created on another Edge node to maintain high availability. If the Edge node that fails is running the active logical router, the original standby logical router becomes the active logical router and a new standby logical router is created. If the Edge node that fails is running the standby logical router, the new standby logical router replaces it.

11 (Optional) Click **Route Advertisement**.

Select one or more of the following:

- **All Static Routes**
- **All NAT IP's**
- **All DNS Forwarder Routes**
- **All LB VIP Routes**
- **All Connected Segments and Service Ports**

- **All LB SNAT IP Routes**
- **All IPSec Local Endpoints**

12 Click **Save**.

13 (Optional) Click **Route Advertisement**.

a In the **Set Route Advertisement Rules** field, click **Set** to add route advertisement rules.

14 (Optional) Click **Additional Settings**.

a For IPv6, you can select or create an **ND Profile** and a **DAD Profile**.

These profiles are used to configure Stateless Address Autoconfiguration (SLAAC) and Duplicate Address Detection (DAD) for IPv6 addresses.

b Select an **Ingress QoS Profile** and an **Egress QoS Profile** for traffic limitations.

These profiles are used to set information rate and burst size for permitted traffic. See Add a Gateway QoS Profile for more information on creating QoS profiles.

If this gateway is linked to a tier-0 gateway, the **Router Links** field shows the link addresses.

15 (Optional) Click **Service Interfaces** and **Set** to configure connections to segments. Required in some topologies such as VLAN-backed segments or one-arm load balancing.

a Click **Add Interface**.

b Enter a name and IP address in CIDR format.

If you configure multicast on this gateway, you must not configure tier-1 addresses as static RP address in the PIM profile.

c Select a segment.

d In the **MTU** field, enter a value between 64 and 9000.

e For **URPF Mode**, you can select **Strict** or **None**.

URPF (Unicast Reverse Path Forwarding) is a security feature.

f Add one or more tags.

g In the **ND Profile** field, select or create a profile.

h Click **Save**.

i (Optional) After you create an interface, you can download the ARP proxies for the gateway by clicking the menu icon (three dots) for the interface and selecting **Download ARP Proxies**.

You can also download the ARP proxy for a specific interface by expanding a gateway and then expanding **Service Interfaces**. Click an interface and click the menu icon (three dots) and select **Download ARP Proxy**.

16 (Optional) Click **Static Routes** and **Set** to configure static routes.

   a Click **Add Static Route**.

   b Enter a name and a network address in the CIDR or IPv6 CIDR format.

   c Click **Set Next Hops** to add next hop information.

   d Click **Save**.

17 (Optional) Click **Multicast** and then the toggle to enable multicast.

   You must select an Edge cluster for this gateway. Also, this gateway must be linked to a tier-0 gateway that has multicast enabled.

18 (Optional) If the tier-1 gateway is connected to a tier-0 gateway, you can download the ARP table of the tier-0 gateway. Do the following:

   a Click the tier-0 gateway from the **Linked Tier-0 Gateway** column.

   b Click the menu icon (3 dots) and select **Download ARP Table**.

   c Select an edge node.

   d Click **Download** to save the .CSV file.

Results

The new gateway is added to the list. For any gateway, you can modify its configurations by clicking the menu icon (3 dots) and select **Edit**. To reconfigure service interfaces or static routes, you do not need to click **Edit**. You only need to click the expand icon (right arrow) for the gateway, expand the **Service Interfaces** or **Static Routes** section, and click the number that is shown. Note that the number must be non-zero. If it is zero, you must edit the gateway.

If NSX Federation is configured, this feature of reconfiguring a gateway by clicking on an entity is applicable to gateways created by the Global Manager (GM) as well. Note that some entities in a GM-created gateway can be modified by the Local Manager, but others cannot. For example, **Static Routes** of a GM-created gateway cannot be modified by the Local Manager. Also, from the Local Manager, you can edit existing **Service Interfaces** of a GM-created gateway but you cannot add an interface.

State Synchronization of Tier-1 Gateways

Connection information of the traffic running on a given tier-1 SR (Service Router) is synchronized to its peer tier-1 SR in active-standby or stateful active-active HA modes. Note that stateful active-active mode is only available starting with NSX 4.0.1.1.

Note State synchronization is not supported for TLS Inspection and IDPS.

In NSX 4.0.0.1, note the following about state synchronization:

- State synchronization is supported for Gateway Firewall, Identity Firewall, NAT, IPSec VPN, DHCP, FQDN analysis, and URL filtering.

- If new sessions were going through a tier-1 SR just before a failover, it might happen that those sessions were not synchronized on the associated tier-1 SR and potentially affect the traffic for those sessions.

Starting with NSX 4.0.1.1, note the following about state synchronization:

- In active-standby mode, state synchronization is supported for Gateway Firewall, Identity Firewall, NAT, IPSec VPN, DHCP, FQDN analysis, and URL filtering.
- In active-active mode, state synchronization is supported for Gateway Firewall, Identity Firewall, NAT, FQDN analysis, and URL filtering. FQDN analysis is only supported with a single sub-cluster. IPSec VPN is not supported.
- If new sessions were going through a tier-1 SR just before a failover, it might happen that those sessions were not synchronized on the associated tier-1 SR and potentially affect the traffic for those sessions.

Segments

4

In NSX, segments are virtual layer 2 domains. A segment was earlier called a logical switch.

There are two types of segments in NSX:

- VLAN-backed segments
- Overlay-backed segments

A VLAN-backed segment is a layer 2 broadcast domain that is implemented as a traditional VLAN in the physical infrastructure. This means that traffic between two VMs on two different hosts but attached to the same VLAN-backed segment is carried over a VLAN between the two hosts. The resulting constraint is that you must provision an appropriate VLAN in the physical infrastructure for those two VMs to communicate at layer 2 over a VLAN-backed segment.

In an overlay-backed segment, traffic between two VMs on different hosts but attached to the same overlay segment have their layer 2 traffic carried by a tunnel between the hosts. NSX instantiates and maintains this IP tunnel without the need for any segment-specific configuration in the physical infrastructure. As a result, the virtual network infrastructure is decoupled from the physical network infrastructure. That is, you can create segments dynamically without any configuration of the physical network infrastructure.

The default number of MAC addresses learned on an overlay-backed segment is 2048. The default MAC limit per segment can be changed through the API field `remote_overlay_mac_limit` in `MacLearningSpec`. For more information see the `MacSwitchingProfile` in the *NSX API Guide*.

This chapter includes the following topics:

- Segment Profiles
- Add a Segment
- Edge Bridging: Extending Overlay Segments to VLAN
- Add a Metadata Proxy Server
- Distributed Port Groups

Segment Profiles

Segment profiles include Layer 2 networking configuration details for segments and segment ports. NSX Manager supports several types of segment profiles.

The following types of segment profiles are available:

- QoS (Quality of Service)
- IP Discovery
- SpoofGuard
- Segment Security
- MAC Management

Note You cannot edit or delete the default segment profiles. If you require alternate settings from what is in the default segment profile you can create a custom segment profile. By default all custom segment profiles except the segment security profile will inherit the settings of the appropriate default segment profile. For example, a custom IP discovery segment profile by default will have the same settings as the default IP discovery segment profile.

Each default or custom segment profile has a unique identifier. You use this identifier to associate the segment profile to a segment or a segment port.

A segment or segment port can be associated with only one segment profile of each type. You cannot have, for example, two QoS segment profiles associated with a segment or segment port.

If you do not associate a segment profile when you create a segment, then the NSX Manager associates a corresponding default system-defined segment profile. The children segment ports inherit the default system-defined segment profile from the parent segment.

When you create or update a segment or segment port you can choose to associate either a default or a custom segment profile. When the segment profile is associated or disassociated from a segment the segment profile for the children segment ports is applied based on the following criteria.

- If the parent segment has a profile associated with it, the child segment port inherits the segment profile from the parent.
- If the parent segment does not have a segment profile associated with it, a default segment profile is assigned to the segment and the segment port inherits that default segment profile.
- If you explicitly associate a custom profile with a segment port, then this custom profile overrides the existing segment profile.

Note If you have associated a custom segment profile with a segment, but want to retain the default segment profile for one of the child segment port, then you must make a copy of the default segment profile and associate it with the specific segment port.

You cannot delete a custom segment profile if it is associated to a segment or a segment port. You can find out whether any segments and segment ports are associated with the custom segment profile by going to the Assigned To section of the Summary view and clicking on the listed segments and segment ports.

Understanding QoS Segment Profile

QoS provides high-quality and dedicated network performance for preferred traffic that requires high bandwidth. The QoS mechanism does this by prioritizing sufficient bandwidth, controlling latency and jitter, and reducing data loss for preferred packets even when there is a network congestion. This level of network service is provided by using the existing network resources efficiently.

For this release, shaping and traffic marking namely, CoS and DSCP is supported. The Layer 2 Class of Service (CoS) allows you to specify priority for data packets when traffic is buffered in the segment due to congestion. The Layer 3 Differentiated Services Code Point (DSCP) detects packets based on their DSCP values. CoS is always applied to the data packet irrespective of the trusted mode.

NSX trusts the DSCP setting applied by a virtual machine or modifying and setting the DSCP value at the segment level. In each case, the DSCP value is propagated to the outer IP header of encapsulated frames. This enables the external physical network to prioritize the traffic based on the DSCP setting on the external header. When DSCP is in the trusted mode, the DSCP value is copied from the inner header. When in the untrusted mode, the DSCP value is not preserved for the inner header.

Note DSCP settings work only on tunneled traffic. These settings do not apply to traffic inside the same hypervisor.

You can use the QoS switching profile to configure the average ingress and egress bandwidth values to set the transmit limit rate. The peak bandwidth rate is used to support burst traffic a segment is allowed to prevent congestion on the northbound network links. These settings do not guarantee the bandwidth but help limit the use of network bandwidth. The actual bandwidth you will observe is determined by the link speed of the port or the values in the switching profile, whichever is lower.

The QoS switching profile settings are applied to the segment and inherited by the child segment port.

Create a QoS Segment Profile

You can define the DSCP value and configure the ingress and egress settings to create a custom QoS switching profile.

Prerequisites

- Familiarize yourself with the QoS switching profile concept. See Understanding QoS Switching Profile.
- Identify the network traffic you want to prioritize.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Select **Networking > Segments > Segment Profiles**.

3 Click **Add Segment Profile** and select **QoS**.

4 Complete the QoS switching profile details.

Option	Description
Name	Name of the profile.
Mode	Select either a **Trusted** or **Untrusted** option from the Mode drop-down menu. When you select the Trusted mode the inner header DSCP value is applied to the outer IP header for IP/IPv6 traffic. For non IP/IPv6 traffic, the outer IP header takes the default value. Trusted mode is supported on an overlay-based logical port. The default value is 0. Untrusted mode is supported on overlay-based and VLAN-based logical port. For the overlay-based logical port, the DSCP value of the outbound IP header is set to the configured value irrespective to the inner packet type for the logical port. For the VLAN-based logical port, the DSCP value of IP/IPv6 packet will be set to the configured value. The DSCP values range for untrusted mode is between 0 to 63. **Note** DSCP settings work only on tunneled traffic. These settings do not apply to traffic inside the same hypervisor.
Priority	Set the DSCP priority value. The priority values range from 0 to 63.
Class of Service	Set the CoS value. CoS is supported on VLAN-based logical port. CoS groups similar types of traffic in the network and each type of traffic is treated as a class with its own level of service priority. The lower priority traffic is slowed down or in some cases dropped to provide better throughput for higher priority traffic. CoS can also be configured for the VLAN ID with zero packet. The CoS values range from 0 to 7, where 0 is the best effort service.
Ingress	Set custom values for the outbound network traffic from the VM to the logical network. You can use the average bandwidth to reduce network congestion. The peak bandwidth rate is used to support burst traffic and the burst size is based on the duration with peak bandwidth. You set burst duration in the burst size setting. You cannot guarantee the bandwidth. However, you can use the Average, Peak, and Burst Size settings to limit network bandwidth. For example, if the average bandwidth is 30 Mbps, peak bandwidth is 60 Mbps, and the allowed duration is 0.1 second, then the burst size is 60 * 1000000 * 0.10/8 = 750000 Bytes. The default value 0 disables rate limiting on the ingress traffic.

Option	Description
Ingress Broadcast	Set custom values for the outbound network traffic from the VM to the logical network based on broadcast. For example, when you set the average bandwidth for a logical switch to 3000 Kbps, peak bandwidth is 6000 Kbps, and the allowed duration is 0.1 second, then the burst size is 6000 * 1000 * 0.10/8 = 75000 Bytes. The default value 0 disables rate limiting on the ingress broadcast traffic.
Egress	Set custom values for the inbound network traffic from the logical network to the VM. The default value 0 disables rate limiting on the egress traffic.

If the ingress, ingress broadcast, and egress options are not configured, the default values are used.

5 Click **Save**.

Understanding IP Discovery Segment Profile

IP Discovery uses DHCP and DHCPv6 snooping, ARP (Address Resolution Protocol) snooping, ND (Neighbor Discovery) snooping, and VM Tools to learn MAC and IP addresses.

Note IP discovery methods for IPv6 are disabled in the default IP discovery segment profile. To enable IP discovery for IPv6 for segments, you must create an IP discovery profile with the IPv6 options enabled and attach the profile to the segments. In addition, make sure that distributed firewall allows IPv6 Neighbor Discovery packets between all workloads (allowed by default).

The discovered MAC and IP addresses are used to achieve ARP/ND suppression, which minimizes traffic between VMs connected to the same segment. The number of IPs in the ARP/ND suppression cache for any given port is determined by the settings in the port's IP Discovery profile. The relevant settings are ARP Binding Limit, ND Snooping Limit, Duplicate IP Detection, ARP ND Binding Limit Timeout, and Trust on First Use (TOFU).

The discovered MAC and IP addresses are also used by the SpoofGuard and distributed firewall (DFW) components. DFW uses the address bindings to determine the IP address of objects in firewall rules.

DHCP/DHCPv6 snooping inspects the DHCP/DHCPv6 packets exchanged between the DHCP/DHCPv6 client and server to learn the IP and MAC addresses.

ARP snooping inspects the outgoing ARP and GARP (gratuitous ARP) packets of a VM to learn the IP and MAC addresses.

VM Tools is software that runs on an ESXi-hosted VM and can provide the VM's configuration information including MAC and IP or IPv6 addresses. This IP discovery method is available for VMs running on ESXi hosts only.

ND snooping is the IPv6 equivalent of ARP snooping. It inspects neighbor solicitation (NS) and neighbor advertisement (NA) messages to learn the IP and MAC addresses.

Duplicate address detection checks whether a newly discovered IP address is already present on the realized binding list for a different port. This check is performed for ports on the same segment. If a duplicate address is detected, the newly discovered address is added to the discovered list, but is not added to the realized binding list. All duplicate IPs have an associated discovery timestamp. If the IP that is on the realized binding list is removed, either by adding it to the ignore binding list or by disabling snooping, the duplicate IP with the oldest timestamp is moved to the realized binding list. The duplicate address information is available through an API call.

By default, the discovery methods ARP snooping and ND snooping operate in a mode called trust on first use (TOFU). In TOFU mode, when an address is discovered and added to the realized bindings list, that binding remains in the realized list forever. TOFU applies to the first 'n' unique <IP, MAC, VLAN> bindings discovered using ARP/ND snooping, where 'n' is the binding limit that you can configure. You can disable TOFU for ARP/ND snooping. The methods will then operate in trust on every use (TOEU) mode. In TOEU mode, when an address is discovered, it is added to the realized bindings list and when it is deleted or expired, it is removed from the realized bindings list. DHCP snooping and VM Tools always operate in TOEU mode.

Note TOFU is not the same as SpoofGuard, and it does not block traffic in the same way as SpoofGuard. For more information, see Understanding SpoofGuard Segment Profile.

For Linux VMs, the ARP flux problem might cause ARP snooping to obtain incorrect information. The problem can be prevented with an ARP filter. For more information, see http://linux-ip.net/html/ether-arp.html#ether-arp-flux.

For each port, NSX Manager maintains an ignore bindings list, which contains IP addresses that cannot be bound to the port. If you navigate to **Networking > Logical Switches > Ports** in **Manager** mode and select a port, you can add discovered bindings to the ignore bindings list. You can also delete an existing discovered or realized binding by copying it to **Ignore Bindings**.

Create an IP Discovery Segment Profile

NSX has several default IP Discovery segment profiles. You can also create additional ones.

IP Discovery is the central infrastructure in NSXwhich determines the set of IP addresses that are associated with a port in the system. IP Discovery policies are applied via the Segment IP Discovery Profile which is configurable from the Policy Manager. It can be associated with a segment, segment port or a group. See Configure IP Discovery Segment Profile on Groups. When a segment or segment port is created, it is initially assigned a Default Segment IP Discovery Profile with a predefined set of policies.

Prerequisites

Familiarize yourself with the IP Discovery segment profile concepts. See Understanding IP Discovery Segment Profile.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Networking > Segments > Segment Profiles**.

3 Click **Add Segment Profile** and select **IP Discovery**.

4 Specify the IP Discovery segment profile details.

Option	Description
Name	Enter a name.
ARP Snooping	For an IPv4 environment. Applicable if VMs have static IP addresses.
ARP Binding Limit	The maximum number of IPv4 IP addresses that can be bound to a port. The minimum value allowed is 1 and the maximum is 256. The default is 1.
ARP ND Binding Limit Timeout	The timeout value, in minutes, for IP addresses in the ARP/ND binding table if TOFU is disabled. If an address times out, a newly discovered address replaces it.
DHCP Snooping	For an IPv4 environment. Applicable if VMs have IPv4 addresses.
DHCP Snooping - IPv6	For an IPv6 environment. Applicable if VMs have IPv6 addresses.
VM Tools	Available for ESXi-hosted VMs only.
VM Tools - IPv6	Available for ESXi-hosted VMs only.
ND Snooping	For an IPv6 environment. Applicable if VMs have static IP addresses.
ND Snooping Limit	The maximum number of IPv6 addresses that can be bound to a port.
Trust on First Use	Applicable to ARP and ND snooping.
Duplicate IP Detection	For all snooping methods and both IPv4 and IPv6 environments.

5 Click **Save**.

Configure IP Discovery Segment Profile on Groups

Configuring IP Discovery segment profiles on a group allows a Security Administrator to configure IP discovery profile parameters, and apply them to group members.

Configuring The following static and dynamic group members are supported:

- Segment
- Segment Port
- VM
- Groups
- Mix of the above

Profiles on groups only apply if the default profile is applied to the segment or segment port:

Custom Group Profile	Custom Profile on Segment (S) and Segment Port(SP)	Effective Profile on Port
Custom	Default (S), Default (SP)	Custom
Custom 1	Default (S), Custom 2 (SP)	Custom 2

Custom Group Profile	Custom Profile on Segment (S) and Segment Port(SP)	Effective Profile on Port
Custom 1	Custom 2 (S), Default (SP)	Custom 2
Custom 1	Custom 2 (S), Custom 3 (SP)	Custom 3

Each time a profile is applied to a group a sequence number is specified. If a member is present in multiple groups, the group with the lower sequence number has higher priority.

Discovery Profile Binding Map API

Method	API	Resource Type
PUT, PATCH, GET, DELETE	`/infra/domains/<domain-id>/groups/<group-id>/discovery-profile-binding-maps/<binding-map-id>`	DiscoveryProfileBindingMap
GET	`/infra/domains/<domain-id>/groups/<group-id>/discovery-profile-binding-maps`	DiscoveryProfileBindingMapListResult

Parameters for DiscoveryProfileBindingMap

Field	Type	Description
profile_path	Policy Path	Required
sequence_number	Integer	Required. Sequence number is used to resolve conflicts when two profiles are applied to the same segment or segment port. The low sequence number has higher precedence.

API for Segments and Ports

Method	API	Resource Type

GET	`/infra/tier-1s/<tier-1-id>/segments/<segment-id>/effective-profiles`
	`/infra/segments/<segment-id>/effective-profiles`
	`/infra/segments/<segment-id>/effective-profiles`
	`/infra/segments/<segment-id>/ports/<port-id>`
		EffectiveProfilesResponse

Example Request

POST https://{{policy-ip}}/policy/api/v1/infra/domains/default/groups/TestGroup/discovery-profile-binding-maps/ipdmap

```
{
    "profile_path" : "/infra/ip-discovery-profiles/ip-discovery-custom-profile-1",
    "sequence_number" :"10"
}
```

Understanding SpoofGuard Segment Profile

SpoofGuard helps prevent a form of malicious attack called "web spoofing" or "phishing." A SpoofGuard policy blocks traffic determined to be spoofed.

SpoofGuard is a tool that is designed to prevent virtual machines in your environment from sending traffic with an IP address it is not authorized to send traffic from. In the instance that a virtual machine's IP address does not match the IP address on the corresponding logical port and segment address binding in SpoofGuard, the virtual machine's vNIC is prevented from accessing the network entirely. SpoofGuard can be configured at the port or segment level. There are several reasons SpoofGuard might be used in your environment:

- Preventing a rogue virtual machine from assuming the IP address of an existing VM.
- Ensuring the IP addresses of virtual machines cannot be altered without intervention – in some environments, it's preferable that virtual machines cannot alter their IP addresses without proper change control review. SpoofGuard facilitates this by ensuring that the virtual machine owner cannot simply alter the IP address and continue working unimpeded.
- Guaranteeing that distributed firewall (DFW) rules will not be inadvertently (or deliberately) bypassed – for DFW rules created utilizing IP sets as sources or destinations, the possibility always exists that a virtual machine could have its IP address forged in the packet header, thereby bypassing the rules in question.

NSX SpoofGuard configuration covers the following:

- MAC SpoofGuard - authenticates MAC address of packet
- IP SpoofGuard - authenticates MAC and IP addresses of packet
- Dynamic Address Resolution Protocol (ARP) inspection, that is, ARP and Gratuitous Address Resolution Protocol (GARP) SpoofGuard and Neighbor Discovery (ND) SpoofGuard validation are all against the MAC source, IP Source and IP-MAC source mapping in the ARP/GARP/ND payload.

At the port level, the allowed MAC/VLAN/IP allow-list is provided through the Address Bindings property of the port. When the virtual machine sends traffic, it is dropped if its IP/MAC/VLAN does not match the IP/MAC/VLAN properties of the port. The port level SpoofGuard deals with traffic authentication, i.e. is the traffic consistent with VIF configuration.

At the segment level, the allowed MAC/VLAN/IP allow-list is provided through the Address Bindings property of the segment. This is typically an allowed IP range/subnet for the segment and the segment level SpoofGuard deals with traffic authorization.

Traffic must be permitted by port level AND segment level SpoofGuard before it will be allowed into segment. Enabling or disabling port and segment level SpoofGuard, can be controlled using the SpoofGuard segment profile.

Create a SpoofGuard Segment Profile

When SpoofGuard is configured, if the IP address of a virtual machine changes, traffic from the virtual machine may be blocked until the corresponding configured port/segment address bindings are updated with the new IP address.

Enable SpoofGuard for the port group(s) containing the guests. When enabled for each network adapter, SpoofGuard inspects packets for the prescribed MAC and its corresponding IP address.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Select **Networking > Segments > Segment Profiles**.
3. Click **Add Segment Profile** and select **Spoof Guard**.
4. Enter a name.
5. To enable port level SpoofGuard, set **Port Bindings** to **Enabled**.
6. Click **Save**.

Understanding Segment Security Segment Profile

Segment security provides stateless Layer2 and Layer 3 security by checking the ingress traffic to the segment and dropping unauthorized packets sent from VMs by matching the IP address, MAC address, and protocols to a set of allowed addresses and protocols. You can use segment security to protect the segment integrity by filtering out malicious attacks from the VMs in the network.

Note that the default segment security profile has the DHCP settings `Server Block` and `Server Block – IPv6` enabled. This means that a segment that uses the default segment security profile will block traffic from a DHCP server to a DHCP client. If you want a segment that allows DHCP server traffic, you must create a custom segment security profile for the segment.

Create a Segment Security Segment Profile

You can create a custom segment security segment profile with MAC destination addresses from the allowed BPDU list and configure rate limiting.

Prerequisites

Familiarize yourself with the segment security segment profile concept. See Understanding Switch Security Switching Profile.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Networking > Segments > Segment Profiles**.

3 Click **Add Segment Profile** and select **Segment Security**.

4 Complete the segment security profile details.

Option	Description
Name	Name of the profile.
BPDU Filter	Toggle the **BPDU Filter** button to enable BPDU filtering. Disabled by default. When the BPDU filter is enabled, all of the traffic to BPDU destination MAC address is blocked. The BPDU filter when enabled also disables STP on the logical switch ports because these ports are not expected to take part in STP.
BPDU Filter Allow List	Click the destination MAC address from the BPDU destination MAC addresses list to allow traffic to the permitted destination. You must enable **BPDU Filter** to be able to select from this list.
DHCP Filter	Toggle the **Server Block** button and **Client Block** button to enable DHCP filtering. Both are disabled by default. DHCP Server Block blocks traffic from a DHCP server to a DHCP client. Note that it does not block traffic from a DHCP server to a DHCP relay agent. DHCP Client Block prevents a VM from acquiring a DHCP IP address by blocking DHCP requests.
DHCPv6 Filter	Toggle the **Server Block - IPv6** button and **Client Block - IPv6** button to enable DHCP filtering. Both are disabled by default. DHCPv6 Server Block blocks traffic from a DHCPv6 server to a DHCPv6 client. Note that it does not block traffic from a DHCP server to a DHCP relay agent. Packets whose UDP source port number is 547 are filtered. DHCPv6 Client Block prevents a VM from acquiring a DHCP IP address by blocking DHCP requests. Packets whose UDP source port number is 546 are filtered.
Block Non-IP Traffic	Toggle the **Block Non-IP Traffic** button to allow only IPv4, IPv6, ARP, and BPDU traffic. The rest of the non-IP traffic is blocked. The permitted IPv4, IPv6, ARP, GARP and BPDU traffic is based on other policies set in address binding and SpoofGuard configuration. By default, this option is disabled to allow non-IP traffic to be handled as regular traffic.
RA Guard	Toggle the **RA Guard** button to filter out ingress IPv6 router advertisements. ICMPv6 type 134 packets are filtered out. This option is enabled by default.
Rate Limits	Set a rate limit for broadcast and multicast traffic. This option is enabled by default. Rate limits can be used to protect the logical switch or VMs from events such as broadcast storms. To avoid any connectivity problems, the minimum rate limit value must be >= 10 pps.

5 Click **Save**.

Understanding MAC Discovery Segment Profile

The MAC discovery segment profile supports two functionalities: MAC learning and MAC address change.

The MAC address change feature allows a VM to change its MAC address. A VM connected to a port can run an administrative command to change the MAC address of its vNIC and still send and receive traffic on that vNIC. This feature is supported on ESXi only. In the default MAC discovery segment profile, this property is enabled.

MAC learning provides network connectivity to deployments where multiple MAC addresses get configured behind one vNIC, for example, in a nested hypervisor deployment where an ESXi VM runs on an ESXi host and multiple VMs run inside the ESXi VM. Without MAC learning, when the vNIC of the ESXi VM connects to a segment port, its MAC address is static. VMs running inside the ESXi VM do not have network connectivity because their packets have different source MAC addresses. With MAC learning, the vSwitch inspects the source MAC address of every packet coming from the vNIC, learns the MAC address and allows the packet to proceed. If a MAC address that is learned is not used for a certain period of time, it is removed. This time period is not configurable. The field **MAC Learning Aging Time** displays the pre-defined value, which is 600.

MAC learning also supports unknown unicast flooding. Normally, when a packet that is received by a port has an unknown destination MAC address, the packet is dropped. With unknown unicast flooding enabled, the port floods unknown unicast traffic to every port on the switch that has MAC learning and unknown unicast flooding enabled. This property is enabled by default, but only if MAC learning is enabled.

The number of MAC addresses that can be learned is configurable. The maximum value is 4096, which is the default. You can also set the policy for when the limit is reached. The options are:

- **Drop** - Packets from an unknown source MAC address are dropped. Packets inbound to this MAC address will be treated as unknown unicast. The port will receive the packets only if it has unknown unicast flooding enabled.
- **Allow** - Packets from an unknown source MAC address are forwarded although the address will not be learned. Packets inbound to this MAC address will be treated as unknown unicast. The port will receive the packets only if it has unknown unicast flooding enabled.

If you enable MAC learning or MAC address change, to improve security, configure SpoofGuard as well.

Create a MAC Discovery Segment Profile

You can create a MAC discovery segment profile to manage MAC addresses.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Networking > Segments > Segment Profiles**.

3 Click **Add Segment Profile** and select **MAC Discovery**.

4 Complete the MAC discovery profile details.

Option	Description
Name	Name of the profile.
MAC Change	Enable or disable the MAC address change feature. The default is disabled.
MAC Learning	Enable or disable the MAC learning feature. The default is disabled.
MAC Limit Policy	Select **Allow** or **Drop**. The default is **Allow**. This option is available if you enable MAC learning
Unknown Unicast Flooding	Enable or disable the unknown unicast flooding feature. The default is enabled. This option is available if you enable MAC learning
MAC Limit	Set the maximum number of MAC addresses. The default is 4096. This option is available if you enable MAC learning
MAC Learning Aging Time	For information only. This option is not configurable. The pre-defined value is 600.

5 Click **Save**.

Add a Segment

You can add two kinds of segments: overlay-backed segments and VLAN-backed segments.

Segments are created as part of a transport zone. There are two types of transport zones: VLAN transport zones and overlay transport zones. A segment created in a VLAN transport zone is a VLAN-backed segment, and a segment created in an overlay transport zone is an overlay-backed segment.

- DHCPv4 relay is supported on a VLAN-backed segment through the Service Interface. Only one DHCP v4 relay or service is supported on a segment.
- For a standalone segment that is not connected to a gateway, only Segment DHCP server is supported.
- For a VLAN segment requiring DHCP server, only Segment DHCP server is supported. Gateway DHCP is not supported on VLAN.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Networking > Segments**.

3 Click **Add Segment**.

4 Enter a name for the segment.

5 Select the type of connectivity for the segment.

Connectivity	Description
None	Select this option when you do not want to connect the segment to any upstream gateway (tier-0 or tier-1). Typically, you want to add a standalone segment in the following scenarios: ■ When you want to create a local testing environment for users that are running workloads on the same subnet. ■ When east-west connectivity with users on the other subnets is not necessary. ■ When north-south connectivity to users outside the data center is not necessary. ■ When you want to configure layer 2 bridging or guest VLAN tagging.
Tier-1	Select this option when you want to connect the segment to a tier-1 gateway.
Tier-0	Select this option when you want to connect the segment to a tier-0 gateway.

Note You can change the connectivity of a gateway-connected segment from one gateway to another gateway (same or different gateway type). In addition, you can change the connectivity of segment from "None" to a tier-0 or tier-1 gateway. The segment connectivity changes are permitted only when the gateways and the connected segments are in the same transport zone. However, if the segment has DHCP configured on it, some restrictions and caveats apply on changing the segment connectivity. For more information, see Scenarios: Impact of Changing Segment Connectivity on DHCP.

6 Enter the Gateway IP address of the subnet in a CIDR format. A segment can contain an IPv4 subnet, or an IPv6 subnet, or both.

- If a segment is not connected to a gateway, subnet is optional.
- If a segment is connected either to a tier-1 or tier-0 gateway, subnet is required.

Subnets of one segment must not overlap with the subnets of other segments in your network. A segment is always associated with a single virtual network identifier (VNI) regardless of whether it is configured with one subnet, two subnets, or no subnet.

7 Select a transport zone, which can be an overlay or a VLAN.

To create a VLAN-backed segment, add the segment in a VLAN transport zone. Similarly, to create an overlay-backed segment, add the segment in an overlay transport zone.

8 (Optional) To configure DHCP on the segment, click **Set DHCP Config**.

For a detailed information about DHCP configuration, see Configure DHCP Service.

9 If the transport zone is of type VLAN, specify a list of VLAN IDs. If the transport zone is of type Overlay, and you want to support layer 2 bridging or guest VLAN tagging, specify a list of VLAN IDs or VLAN ranges

10 (Optional) Select an uplink teaming policy for the segment.

This drop-down menu displays the named teaming policies, if you have added them in the VLAN transport zone. If no uplink teaming policy is selected, the default teaming policy is used.

- Named teaming policies are not applicable to overlay segments. Overlay segments always follow the default teaming policy.
- For VLAN-backed segments, you have the flexibility to override the default teaming policy with a selected named teaming policy. This capability is provided so that you can steer the infrastructure traffic from the host to specific VLAN segments in the VLAN transport zone. Before adding the VLAN segment, ensure that the named teaming policy names are added in the VLAN transport zone.

11 (Optional) Enter the fully qualified domain name.

DHCPv4 server and DHCPv4 static bindings on the segment automatically inherit the domain name from the segment configuration as the Domain Name option.

12 If you want to use Layer 2 VPN to extend the segment, click the **L2 VPN** text box and select an L2 VPN server or client session.

You can select more than one.

13 In **VPN Tunnel ID**, enter a unique value that is used to identify the segment.

14 In the **Metadata Proxy** field, select a metadata proxy from the drop-down list, or click the menu icon (3 dots) to create one.

15 Click **Save**.

16 To add segment ports, click **Yes** when prompted if you want to continue configuring the segment.

a Click **Set** from the **Ports / Interfaces** column.

b Click **Add Segment Port**.

c Enter a port name.

d For **ID**, enter the VIF UUID of the VM or server that connects to this port.

e Select a type: **Child**, or **Static**.

Leave this text box blank except for use cases such as containers or VMware HCX. If this port is for a container in a VM, select **Child**. If this port is for a bare metal container or server, select **Static**.

f Enter a context ID.

Enter the parent VIF ID if **Type** is **Child**, or transport node ID if **Type** is **Static**.

g Enter a traffic tag.

Enter the VLAN ID in container and other use cases.

h Select an address allocation method: **IP Pool**, **MAC Pool**, **Both**, or **None**.

i Specify tags.

j Apply address binding by specifying the IP (IPv4 address, IPv4 subnet, IPv6 address, or IPv6 subnet) and MAC address of the logical port to which you want to apply address binding. For example, for IPv6, 2001::/64 is an IPv6 subnet, 2001::1 is a host IP, whereas 2001::1/64 is an invalid input. You can also specify a VLAN ID.

Manual address bindings, if specified, override the auto discovered address bindings.

k Select segment profiles for this port.

17 To select segment profiles, click **Segment Profiles**.

18 Click **Save**.

19 (Optional) You can click the menu icon (3 dots) for the following options to save specific information about the segment in a CSV file:

- **Download MAC Table**: Select the source which can be the Central Control Plane or a specific transport node for the associated MAC addresses.
- **Download VTEP Table**: Select the source which can be the Central Control Plane or a specific transport node for the associated VTEPs.
- **Download ARP Table**: Select the edge node to save the ARP table.

  Note This option is only available if the segment is connected to a gateway.

- **Download ARP Proxy**: Save the aggregate of the ARP proxy for the segment.

  Note This option is only available if the segment is connected to a gateway.

20 (Optional) You can view more information about the segment by expanding the segment and clicking the following options on the right:

- **View Statistics**: Contains the following tabs:
  - **Local Port**: Displays the traffic details for the local port.
  - **Interface Statistics**: Displays the data details for specific edge nodes.

    Note This tab is only available if the segment is connected to the gateway.

  - **DHCP Statistics**: Displays the DHCP server packet counts and DHCP pool usage statistics. This tab is available only when you have configured Segment DHCP server on the segment.

    If you have configured both DHCPv4 and DHCPv6 servers on a segment, the **DHCP Statistics** tab will display only the DHCPv4 packet counts and the DHCPv4 pool usage statistics. DHCPv6 packet counts and DHCPv6 pool usage statistics are currently not supported.
- **View Related Groups**: Displays the groups associated with the segment.

- **View DAD Status**: Displays Duplicate Address Detection (DAD) status for the segment.

  Note This tab is only available if the segment is connected to the gateway.

Results

The new segment is added to the list. For any segment, you can modify its configurations by clicking the menu icon (3 dots) and select **Edit**. To reconfigure ports, you do not need to click **Edit**. You only need to click the expand icon (right arrow) for the segment and click the number in the **Ports** column. Note that the number must be non-zero. If it is zero, you must edit the segment.

If NSX Federation is configured, this feature of reconfiguring a segment by clicking on an entity is applicable to segments created by the Global Manager (GM) as well. Note that from the Local Manager, you can create ports for a GM-created segment because you cannot create segment ports from the Global Manager.

Edge Bridging: Extending Overlay Segments to VLAN

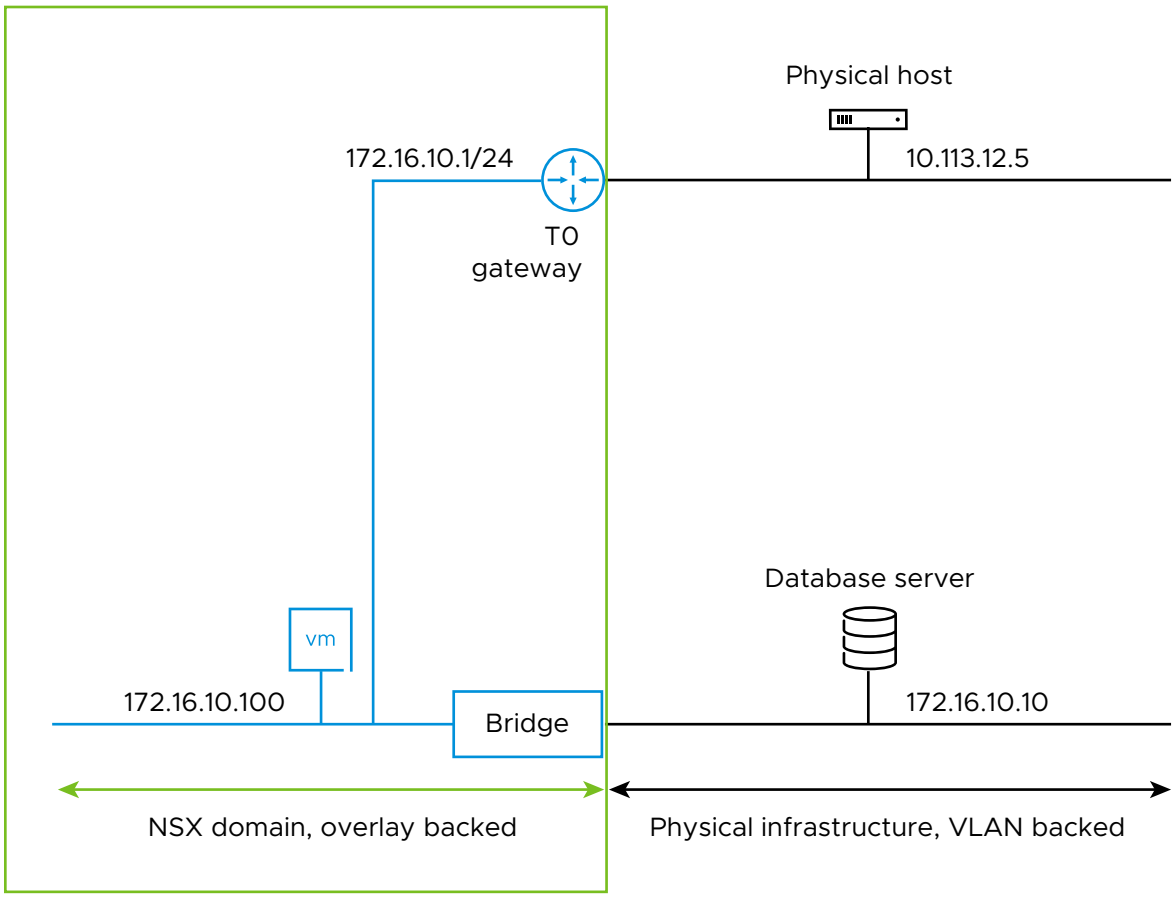
Workloads attached to overlay segments typically communicate at layer 3 with physical devices outside of the NSX domain, through tier-0 gateways instantiated on NSX Edge. However, there are some scenarios where layer 2 connectivity is required between virtual machines in NSX and physical devices.

Some examples are:

- Migration from physical to virtual, or virtual to virtual.
- Integration of a physical appliance that provides services to a segment, like an external load balancer.
- Connection to a database server that requires layer 2 adjacency to its virtual machine clients.

For that purpose, on the top of the gateway service, NSX Edge can also run a bridge service. The following diagram represents those two options: the virtual machine in the bottom left corner has layer 3 connectivity through a gateway to the physical host, and layer 2 connectivity through a bridge to the database server. It is possible to both route and bridge a segment. In fact, it is possible to use the tier 0 gateway in this diagram as a default gateway for the database server.

Figure 4-1. NSX VM Bridge and Gateway Communication



The NSX Edge bridge, like the gateway, is supported for long term deployments, even if it is often used as a temporary solution during migrations.

The bridge functionality extends an overlay segment into a VLAN, identified by a VLAN ID on an uplink of the NSX Edge where the bridge is running. Typically, two redundant active and standby bridges get deployed on separate edges as part of the same edge cluster. There is no active/active redundancy possible. Setting up the bridge functionality involves the following configuration steps:

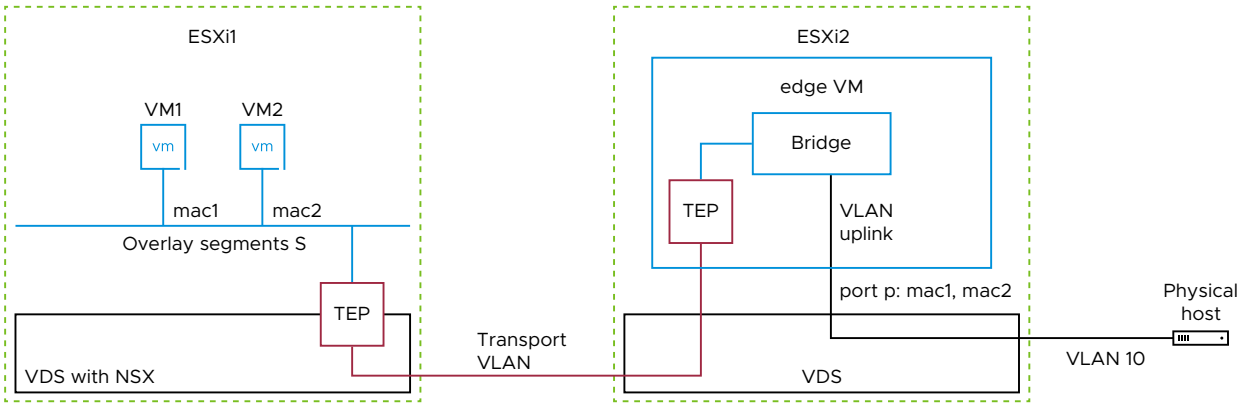
- Make sure that the NSX Edge is suitable for hosting the bridge service. The bridge adds a few constraints to the deployment of an edge in a VM form factor.
- Identify the NSX Edges that run the bridge service. A bridge profile statically designates the edge responsible for running the active bridge and optionally designates a second edge hosting the standby bridge.
- Lastly associate an overlay segment to a VLAN ID or IDs and a bridge profile. This results in the creation of the appropriate active/standby bridges on the edges specified in the bridge profile, that extend at layer 2 the overlay segment to the VLAN or VLANs identified by the VLAN IDs.

Configure an Edge VM for Bridging

There are no specific constraints to configure bridging on a bare metal edge. However, if you are planning to run a bridge on an NSX Edge VM, use this section to understand the specific configuration to perform in the vSphere infrastructure.

As an example, our scenario includes two virtual machines, VM1 and VM2, on transport node ESXi 1 attached to an overlay segment S. The VMs can communicate at layer 2 with the physical host on the right side of the diagram thanks to a bridge instantiated on the edge VM running on ESXi host 2. The TEP (tunnel end point) on ESXi 1 encapsulates the traffic from VM1/VM2 and forwards it to the TEP of the edge VM. Then the bridge unencapsulates the traffic and sends it tagged with VLAN ID 10 on its VLAN uplink. Then the traffic gets switched to the physical host.

Figure 4-2. Edge VM Bridging



As you can see on the diagram, the VLAN uplink of the bridge is linked to port p, which is attached to distributed portgroup dvpg1. This port p is injecting into dvpg1 traffic with the source mac addresses mac1 and mac2 of virtual machined VM1 and VM2. Port p must also accept traffic with destination mac addresses mac1 and mac2, so that the physical host can reach VM1 and VM2. When a bridge is running on the edge VM, the port of this edge VM behaves in a non-standard way as far as the vSphere switching infrastructure is concerned. That means that dvpg1 will need some additional configuration to accommodate the edge VM. The following section lists the different options based on your environment.

Option 1: Edge VM is on a VSS portgroup

This option is for when the Edge VM is connected to a VSS (vSphere Standard Switch). You must enable promiscuous mode and forged transmit.

- Set promiscuous mode on the portgroup.
- Allow forged transmit on the portgroup.
- Run the following command to enable reverse filter on the ESXi host where the Edge VM is running:

```
esxcli system settings advanced set -o /Net/ReversePathFwdCheckPromisc -i 1
```

Then disable and enable promiscuous mode on the portgroup with the following steps:

- Edit the portgroup's settings.
- Disable promiscuous mode and save the settings.
- Edit the portgroup's settings again.
- Enable promiscuous mode and save the settings.

- Do not have other port groups in promiscuous mode on the same host sharing the same set of VLANs.
- Avoid running other VMs attached to the portgroup in promiscuous mode on the same host, as the traffic gets replicated to all those VMs and affect performance.

Option 2a: Edge VM is on a VDS 6.6.0 (or later) portgroup

This option is for when the Edge VM is connected to a VDS (vSphere Distributed Switch). You must be running ESXi 6.7 or later, and VDS 6.6.0 or later.

- Enable MAC learning with the option “allow unicast flooding” on the distributed portgroup.

  Starting with vSphere 8.0, you can enable the Mac Learning UI option in the distributed portgroup configuration. For previous releases, you need to use the VIM API `DVSMacLearningPolicy` and setting `allowUnicastFlooding` to `true`. For an example of how to enable MAC learning, see https://williamlam.com/2018/04/native-mac-learning-in-vsphere-6-7-removes-the-need-for-promiscuous-mode-for-nested-esxi.html.

Option 2b: Edge VM is on a VDS 6.5.0 (or later) portgroup

This option is for when the Edge VM is connected to a VDS (vSphere Distributed Switch). You enable promiscuous mode and forged transmit.

- Set promiscuous mode on the distributed portgroup.
- Allow forged transmit on the distributed portgroup.
- Run the following command to enable reverse filter on the ESXi host where the Edge VM is running:

  ```
  esxcli system settings advanced set -o /Net/ReversePathFwdCheckPromisc -i 1
  ```

  Then disable and enable promiscuous mode on the distributed portgroup with the following steps:

  - Edit the distributed portgroup's settings.
  - Disable promiscuous mode and save the settings.
  - Edit the distributed portgroup's settings again.
  - Enable promiscuous mode and save the settings.

- Do not have other distributed port groups in promiscuous mode on the same host sharing the same set of VLANs.

- Avoid running other VMs attached to the distributed portgroup in promiscuous mode on the same host, as the traffic gets replicated to all those VMs and affects performance.

Option 3: Edge VM is connected to an NSX segment

If the Edge is deployed on a host with NSX installed, it can connect to a VLAN segment and use MAC Learning, which is the preferred configuration option.

- Create a new MAC Discovery segment profile by navigating to **Networking > Segments > Profiles**.
  - Click **Add Segment Profile > MAC Discovery**.
  - Enable **MAC Learning**. This will also enable **Unknown Unicast Flooding**. Keep the flooding option enabled for bridging to work in all scenarios.
  - Click **Save**.
- Edit the segment used by the Edge by navigating to **Networking > Segments**.
  - Click the menu icon (3 dots) and select **Edit**.
  - Expand the **Segment Profiles** section, then set the **MAC Discovery** profile to the one created above.

Note If you bridge a segment to VLAN 0 and you use a distributed router on this segment, the gateway might not route VLAN 0 traffic when using MAC learning. In this scenario, avoid option 3. Avoid option 2a if the edge VM is attached to the distributed portgroup of a VDS prepared for NSX for vSphere.

Create an Edge Bridge Profile

The edge bridge profile is a template for instantiating bridges. In the template, you define a primary edge, an optional backup edge from the same edge cluster as the primary, and a failover mode, preemptive or non-preemptive.

The preference is to use the primary edge for running the active bridge, the bridge forwarding traffic between overlay segment and VLAN. The standby bridge, that typically runs on the backup edge, does not forward any traffic.

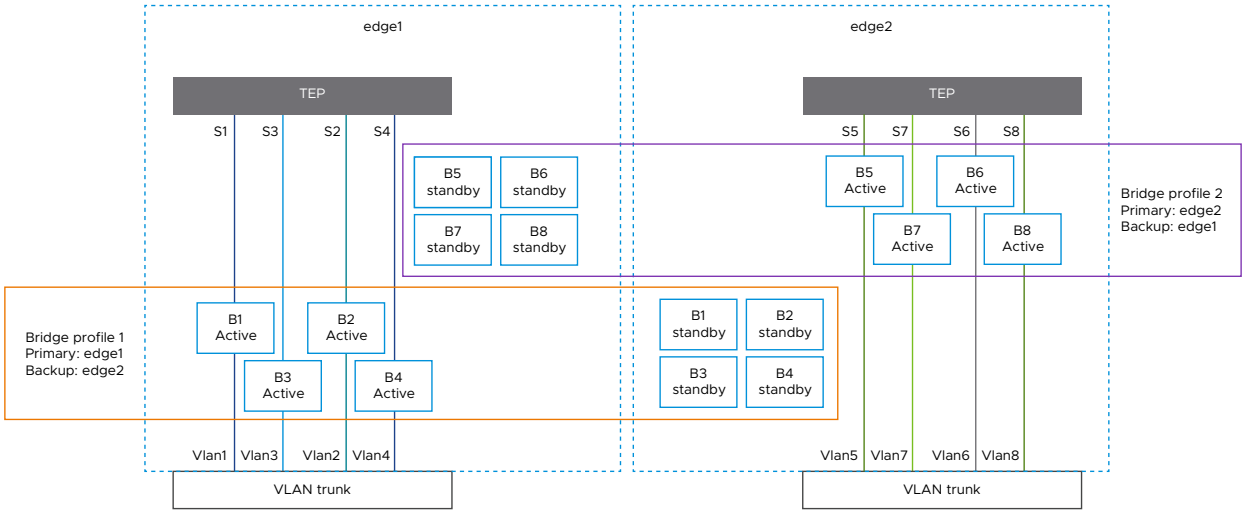
You can instantiate multiple bridges from the same bridge profile. As a result, in most cases, few edge bridge profiles are required. For example, if you plan to use two edges (edge1 and edge2) for bridging, you might want to create two edge bridge profiles:

- Profile 1 with edge1 as primary and edge2 as backup
- Profile 2 with edge2 as primary and edge1 as backup

You can then create an arbitrary number of bridges using edge1 as primary (respectively backup), by associating them to the profile 1 (respectively profile 2). Those two profiles are enough to load share the bridged traffic between the two edges, on a per segment basis. The Few Bridge Profiles for Many Bridges diagram represents an example of bridging eight segments across two edges, using two edge bridge profiles.

This diagram shows bridge overlay segments S1 to VLAN 1, segment S2 to VLAN 2, and so on. Segments S1 to S4 are using bridge profile 1, resulting in active bridges on edge1, standby on edge2. Segment S5 to S8 are using bridge profile 2, leading to active bridges on edge2, standby on edge1. This diagram shows load sharing of the bridging functionality, on a per segment basis.

Figure 4-3. Few Bridge Profiles for Many Bridges



Depending on the availability of the edges and the failover mode selected for the bridge profiles, the active bridges might be running on the backup edges.

When both edges in the bridge profile are available, the active bridge is typically running on the primary edge. If the active bridge or the primary edge fails, the standby bridge on the backup edge takes over the active role and starts forwarding traffic between overlay segment and VLAN.

A bridge switchover, moving the active bridge to a different edge, is an operation that results in traffic loss. The bridge that is becoming active synchronizes the mac addresses that were learned on the previously active bridge and starts flooding RARP packets, using those mac addresses as source mac addresses. This mechanism is necessary to update the mac address tables of the physical infrastructure.

For example, what if a failure occurs on the primary edge and the bridge running on the backup edge is already active? In preemptive mode, when the failure is recovered on the primary edge, a bridge switchover is triggered and the bridge on the primary edge becomes active again.

The benefit of the preemptive mode is that the system is attempting to forward the bridged traffic along a deterministic path. If you take the example of the Few Bridge Profiles for Many Bridges figure, with a preemptive mode, you are sure that the bridge traffic gets distributed on a per segment basis as soon as both edges are available, thus providing more bandwidth.

The drawback of the preemptive mode is that there is a disruptive bridge convergence when the bridge on the primary edge recovers and becomes active again.

In non-preemptive mode, the bridge on the primary edge recovers from failure as a standby bridge. The benefit of this mode is that there is no additional traffic disruption when the primary recovers. The preemptive mode is the best option in terms of bandwidth, thanks to its load sharing. The drawback of the non-preemptive mode is that bridge traffic flow is non-deterministic and can be sub-optimal. In the example shown in the Few Bridge Profiles for Many Bridges figure, after a failed edge recovers, the bridge traffic still flows through a unique edge, with no load sharing.

You can manually trigger a bridge switchover. To manually trigger a bridge switchover from the CLI of the edge currently hosting the standby bridge, enter: `set bridge <uuid> state active`.

Use this command only in non-preemptive mode. If you use it in preemptive mode, it returns an error.

For more information on set or get bridge commands, see the *NSX Command-Line Interface Reference*.

Esnure you verify that you have an NSX Edge cluster with two NSX Edge transport nodes.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Select **Networking > Segments > Profiles**.
3. Click **Edge Bridge Profiles**.
4. Click **Add Edge Bridge Profile**.
5. Enter a name for the Edge bridge profile and optionally a description.
6. Select an NSX Edge cluster.
7. Select a primary node.
8. Select a backup node.
9. Select a failover mode.

   The options are **Preemptive** and **Non-Preemptive**.
10. Click **Save**.

What to do next

Create a bridge-backed segment. See Extend an Overlay Segment to a VLAN or a Range of VLANs.

Extend an Overlay Segment to a VLAN or a Range of VLANs

After you have identified the edges on which you want the bridging functionality to be performed and created the appropriate edge bridge profile, the final step is to edit the segment configuration and specify the edge bridge profile to which you want to associate with the segment and the VLAN ID or range of VLAN IDs to which to bridge your segment. This step instantiates one or two bridges on the edges identified in the edge bridge profile.

When you configure a bridge with a single VLAN ID, a frame received on the overlay segment by the bridge gets decapsulated and forwarded on the VLAN uplink of the bridge with an added 802.1Q tag corresponding to this VLAN ID.

When you create the bridge specifying a VLAN ID range, you must configure the overlay segment being bridged for Guest VLAN Tagging (GVT). This means that the encapsulated frames already carry an 802.1Q tag. When the bridge receives an encapsulated frame carrying a VLAN tag on its overlay interface, it first checks that VLAN ID in the tag belongs to the VLAN range configured for the bridge. After confirmation, it forwards the frame on the VLAN uplink of the bridge carrying the original 802.1Q tag that was received on the overlay. Otherwise, it drops the frame.

Note If needed, you can configure multiple bridges on the same segment, but:

- The same segment cannot be bridged twice on the same edge.
- The bridge does not have any loop detection or prevention. If you configure multiple bridges to the same bridging domain on the VLAN side, it results in a permanent bridging loop.

Configuring a Bridge-Backed Segment

Prerequisites

- You have identified an overlay segment you want to bridge.
- You have an edge bridge profile specifying one or two edges attached to the overlay transport zone of your segment.
- If you are using edge VMs, you have checked the configuration requirements in Configure an Edge VM for Bridging.

Procedure

1. From a browser, log in with admin privileges to an NSX Manager or a Global Manager at `https://<nsx-mgr-or-global-mgr-ip-address>`.
2. Select **Networking > Segments**.
3. Click the menu icon (three dots) of the overlay segment that you want to configure layer 2 bridging on and select **Edit**.
4. Expand **Additional Settings** and in the **Edge Bridges** field, click **Set**.
5. Click **Add Edge Bridge**.

6 Select an Edge bridge profile.

7 Select a VLAN transport zone to identify the VLAN uplinks used by the bridge.

8 Enter a VLAN ID or a VLAN ID range (specify VLAN ranges and not individual VLANs).

9 (Optional) Select a teaming policy.

If there are multiple VLAN uplinks on the edge NVDS attached to the VLAN transport zone selected in the previous steps, use a failover order named teaming policy to identify the exact uplink on which VLAN bridged traffic gets forwarded. The uplinks of a VM edge do not fail, so the teaming policy only needs a single uplink. If you do not enter a specific teaming policy and there are multiple VLAN uplinks, the first one configured on the edge NVDS is used.

10 Click **Add**.

11 Click **Apply**.

Add a Metadata Proxy Server

A metadata proxy server enables VMs to retrieve metadata from an OpenStack Nova API server.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Networking > Segments > Metadata Proxies**.

3 Click **Add Metadata Proxy**.

4 Enter a name for the metadata proxy server.

5 In the **Server Address** field, enter the URL and port for the Nova server.

The valid port range is 3000 - 9000.

6 Select an Edge cluster.

7 (Optional) Select Edge nodes.

If you select any Edge node, you cannot enable **Standby Relocation** in the next step.

8 (Optional) Enable **Standby Relocation**.

Standby relocation means that if the Edge node running the metadata proxy fails, the metadata proxy will run on a standby Edge node. You can only enable standby relocation if you do not select any Edge node.

9 In the **Shared Signature Secret** field, enter the secret that the metadata proxy will use to access the Nova server.

10 (Optional) Select a certificate for encrypted communication with the Nova server.

11 (Optional) Select a cryptographic protocol.

The options are TLSv1, TLSv1.1, and TLSv1.2. TLSv1.1 and TLSv1.2 are supported by default.

Distributed Port Groups

A distributed port group specifies port configuration options for each member port on a vSphere distributed switch. Distributed port groups define how a connection is made to a network.

Distributed Port Group Creation in NSX

When you install Distributed Security to a vSphere Distributed Switch (VDS), the Distributed Virtual port groups (DVPG) and DVports of the VDS are discovered and objects are automatically created to represent them in NSX. For details, see Install Distributed Security for vSphere Distributed Switch.

Important The objects created in NSX Manager for the DVPGs are called *distributed port groups*. They are not called *segments* in the UI.

The objects created in NSX Manager for the DVports are called *distributed ports*. They are not called *segment ports* in the UI.

Also, the following events occur during the Distributed Security installation:

- The VLAN tags for the DVPG are automatically discovered and shown in NSX Manager. The VLAN tags can only be edited in VMware vCenter.
- The default segment profiles are applied to the distributed port groups. You can later switch them to custom profiles.
- Only connected DVports from VMware vCenter are discovered by NSX. Free DVports are not discovered.

After the Distributed Security installation, you can view these objects by navigating to **Networking > Segments**, and then selecting the **Distributed Port Groups** tab.

The distributed port group and distributed port objects are kept in sync between VMware vCenter and NSX. This means that if DVPGs or DVports are created or removed in VMware vCenter, then those changes are automatically made to the respective distributed port groups or distributed ports in NSX Manager. If changes are made in VMware vCenter while connectivity is lost between VMware vCenter and NSX, those changes are automatically processed and reflected in NSX Manager when connectivity is restored.

Available Actions for Distributed Port Groups and Distributed Ports

You can perform the following actions for distributed port groups and distributed ports in NSX Manager:

Object	Available Actions
Distributed port groups	■ Apply SpoofGuard. ■ Apply IP Discovery. ■ Apply switch security profile. ■ Add and remove tags which allows the distributed port group to be added to dynamic NSGroups. ■ Add and remove the distributed port group from static NSGroups.
Distributed ports	■ Add and remove tags which allows distributed ports to be added to dynamic NSGroups. ■ Add and remove distributed ports from static NSGroups. ■ Manage address bindings.

DHCP

5

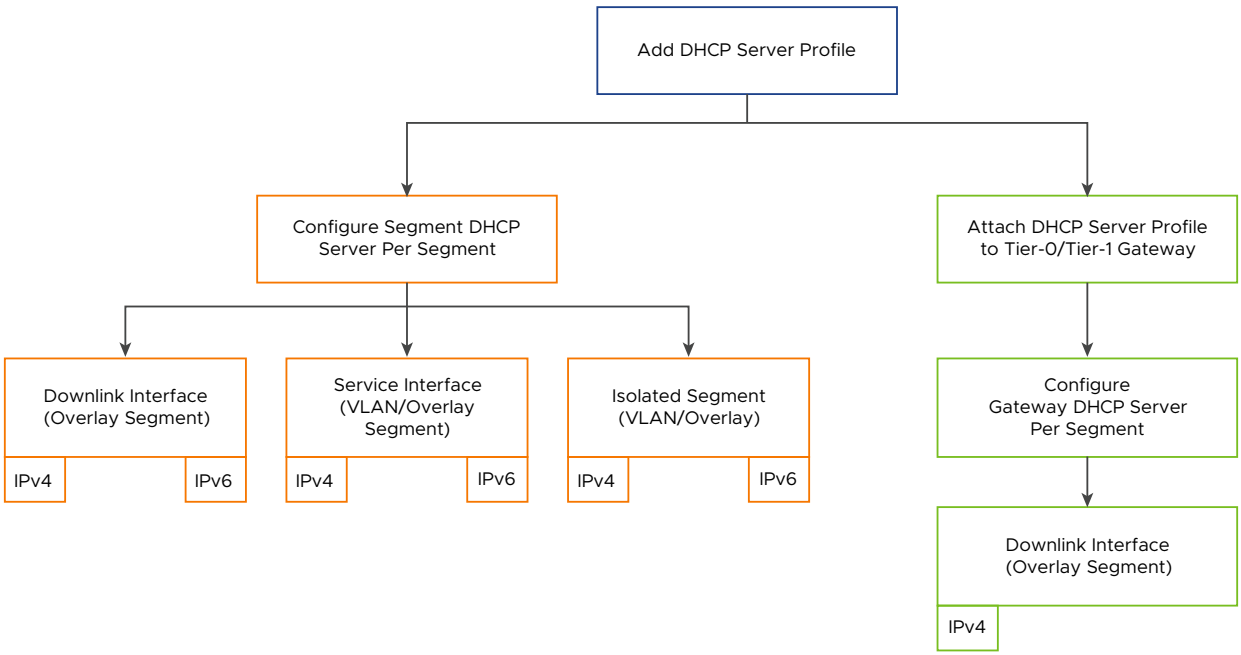
You can configure DHCP service on each segment regardless of whether the segment is connected to a gateway. Both DHCP for IPv4 (DHCPv4) and DHCP for IPv6 (DHCPv6) servers are supported.

NSX supports the following types of DHCP configuration on a segment:

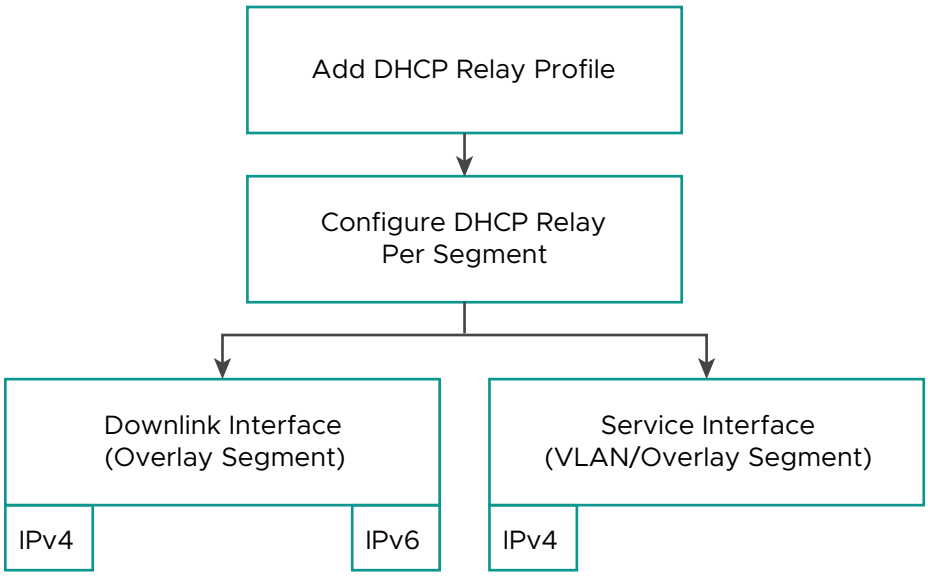
- Segment DHCP server (earlier known as Local DHCP server in NSX 3.x releases)
- Gateway DHCP server (supported only for IPv4 subnets in a segment)
- DHCP Relay

High-level Overview of Configuration

The following figure shows the high-level overview of DHCP server configuration in NSX.



The following figure shows the high-level overview of DHCP Relay configuration in NSX.



Supported DHCP Configuration Types

The following figure shows an example of the various scenarios for Segment DHCP server, Gateway DHCP server, and DHCP Relay in an NSX network.

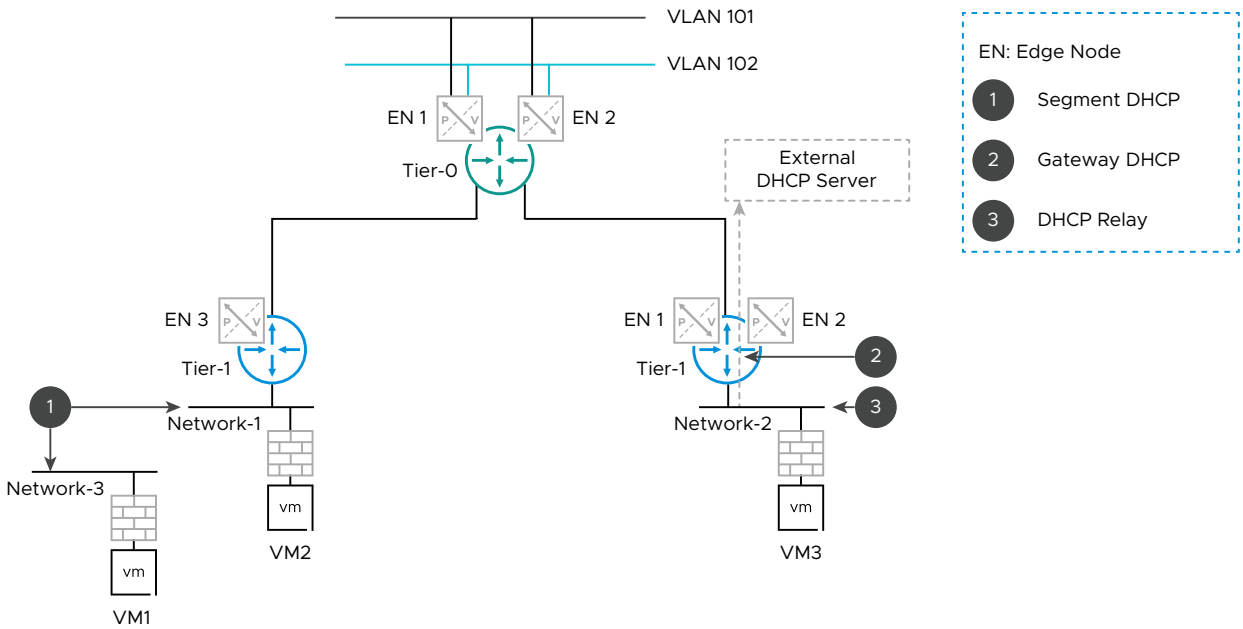


Table 5-1. Types of DHCP Configuration in NSX

DHCP Configuration Type	Description
Segment DHCP server	Select this option to create a Segment DHCP server that has an IP address on the segment. A Segment DHCP server provides a dynamic IP assignment service only to the VMs that are attached to the segment. The DHCP server IP address must belong to the subnet that is configured on the segment. Also, the server IP address must be different from the Gateway IP address of the segment. Segment DHCP server is local to the segment and not available to the other segments in the network. You can configure all DHCP settings, including DHCP ranges, DHCP options, and static bindings on the segment. For isolated segments, which are not connected to a gateway, Segment DHCP server configuration type is selected by default. You can configure DHCPv6 in the IPv6 subnet of a segment with a Segment DHCP server.
DHCP Relay	Select this option to relay the DHCP client requests to the external DHCP servers. The external DHCP servers can be in any subnet, outside the SDDC, or in the physical network. DHCP Relay service is local to the segment and is not available to the other segments in the network. When you use a DHCP Relay on a segment, you cannot configure DHCP settings, DHCP options, and static bindings on the segment.
Gateway DHCP server	Gateway DHCP server is analogous to a central DHCP server that dynamically assigns IP and other network configuration to the VMs on all the segments that are connected to the gateway and using Gateway DHCP server. By default, segments that are connected to a tier-1 or tier-0 gateway use Segment DHCP server. If needed, you can choose to configure a Gateway DHCP server or a DHCP Relay on the segment. To configure Gateway DHCP server on a segment, a DHCP server profile must be attached to the gateway. If the IPv4 subnet of a segment uses a Gateway DHCP server, you cannot configure DHCPv6 in the IPv6 subnet of the same segment because Gateway DHCPv6 server is not supported. In this case, the IPv6 subnet cannot support any DHCPv6 server configuration, including the IPv6 static bindings.

This chapter includes the following topics:

- Configure DHCP Service
- Attach a DHCP Profile to a Tier-0 or Tier-1 Gateway
- View Gateway DHCP Statistics
- View Segment DHCP Statistics
- Scenarios: Selection of Edge Cluster for DHCP Service
- Scenarios: Impact of Changing Segment Connectivity on DHCP

Configure DHCP Service

You can configure DHCP service on each segment regardless of whether the segment is connected to a gateway. Both DHCP for IPv4 (DHCPv4) and DHCP for IPv6 (DHCPv6) servers are supported.

DHCP Configuration Settings: Reference

Use this reference documentation to understand the various considerations that you must keep in mind while configuring the DHCP service and to obtain a detailed guidance about the configuration settings on the **Set DHCP Config** page.

The following note mentions the DHCP configuration types that are supported or not supported based on how overlay or VLAN segments are connected:

Note

- On an isolated segment that is not connected to a gateway, only Segment DHCP server is supported.
- Segments that are configured with an IPv6 subnet can have either a Segment DHCPv6 server or a DHCPv6 relay. Gateway DHCPv6 server is not supported.
- If a segment contains an IPv4 subnet and an IPv6 subnet, you can configure both DHCPv4 and DHCPv6 servers on the segment.
- DHCPv4 relay is supported on a VLAN segment through the service interface of a tier-0 or tier-1 gateway. Only one DHCPv4 relay service is supported on a VLAN segment.
- For a VLAN segment requiring a DHCP server, only Segment DHCP server is supported. Gateway DHCP server is not supported on a VLAN segment.

When a Segment DHCP server or DHCP Relay is configured on a VLAN segment, **Forged Transmits** option on the ESXi host where the Edge VM is installed must be set to **Accept**. Set this option in the VDS or VSS configuration.

In vSphere 7.0 or later, the **Forged Transmits** option is set to **Reject**, by default. To enable this option on the host, do the following steps:

1 Log in to the vSphere Client UI with **admin** privileges.

2 Go to **Hosts and Clusters** and click a host in the cluster.

3 Navigate to **Configure > Networking > Virtual Switches**, and then click **Edit**.

4 On the **Edit Settings** window, click **Security**. From the **Forged Transmits** drop-down menu, select **Accept**.

To learn about Forged Transmits, see the *vSphere Security* documentation.

Important After a segment has DHCP service configured on it, some restrictions and caveats apply on changing the connectivity of the segment. For more information, see Scenarios: Impact of Changing Segment Connectivity on DHCP.

The following sections provide guidance about the configuration settings on the **Set DHCP Config** page.

DHCP Type

- When a segment is connected to a gateway, Segment DHCP server is selected by default. If needed, you can select Gateway DHCP server or DHCP Relay from the drop-down menu.
- If you select the DHCP type as Gateway DHCP server, the DHCP profile that is attached to the gateway is autoselected. The name and server IP address are fetched automatically from that DHCP profile and displayed in a read-only mode.
- For an isolated segment, which is not connected to a gateway, Segment DHCP server is selected by default.

DHCP Profile

- When you are configuring a Segment DHCP server or a DHCP Relay on the segment, you must select a DHCP profile from the drop-down menu. If no profiles are available in the **DHCP Profile** drop-down menu, click  and create a DHCP profile. After the profile is created, it is automatically attached to the segment.
- When a segment is using a Gateway DHCP server, ensure that an edge cluster is selected either in the gateway, or DHCP server profile, or both. If an edge cluster is unavailable in either the profile or the gateway, an error message is displayed when you save the segment.
- When a segment is using a Segment DHCP server, ensure that the DHCP server profile contains an edge cluster. If an edge cluster is unavailable in the profile, an error message is displayed when you save the segment.

IPv4 Server or IPv6 Server Settings

This section explains the configuration settings in the **IPv4 Server** tab page and the **IPv6 Server** tab page.

DHCP Server Address

- If you are configuring a Segment DHCP server, server IP address is required. A maximum of two server IP addresses are supported. One IPv4 address and one IPv6 address. For an IPv4 address, the prefix length must be <= 30, and for an IPv6 address, the prefix length must be <= 126. The server IP addresses must belong to the subnets that you have specified in this segment. The DHCP server IP address must not overlap with the IP addresses in the DHCP ranges and DHCP static binding. The DHCP server profile might contain server IP addresses, but these IP addresses are ignored when you configure a Segment DHCP server on the segment.
- After a Segment DHCP server is created, you can edit the server IP addresses on the **Set DHCP Config** page. However, the new IP addresses must belong to the same subnet that is configured in the segment.

- If you are configuring a Gateway DHCP server, the **DHCP Server Address** text box is not editable. The server IP addresses are fetched automatically from the DHCP profile that is attached to the connected gateway.
- The Gateway DHCP server IP addresses in the DHCP server profile can be different from the subnet that is configured in the segment. In this case, the Gateway DHCP server connects with the IPv4 subnet of the segment through an internal relay service, which is autocreated when the Gateway DHCP server is created. The internal relay service uses any one IP address from the subnet of the Gateway DHCP server IP address.
- The IP address used by the internal relay service acts as the default gateway on the Gateway DHCP server to communicate with the IPv4 subnet of the segment.
- After a Gateway DHCP server is created, you can edit the server IP addresses in the DHCP profile of the gateway. However, you cannot change the DHCP profile that is attached to the gateway. When a DHCP server profile is used in your network, preferably avoid editing the server IP addresses in the DHCP server profile. It might cause a failure while renewing or releasing the IP addresses that are leased to the DHCP clients.

DHCP Ranges

- IP ranges, CIDR subnet, and IP addresses are allowed. IPv4 addresses must be in a CIDR /32 format, and IPv6 addresses must be in a CIDR /128 format. You can also enter an IP address as a range by entering the same IP address in the start and the end of the range. For example, 172.16.10.10-172.16.10.10.
- IP addresses in the DHCP ranges must belong to the subnet that is configured on the segment. That is, DHCP ranges cannot contain IP addresses from multiple subnets.
- IP ranges must not overlap with the DHCP server IP address and the DHCP static binding IP addresses.
- IP ranges in the DHCP IP pool must not overlap each other.
- Number of IP addresses in any DHCP range must not exceed 65536.
- The following types of IPv6 addresses are not permitted in DHCP for IPv6 ranges:
  - Link Local Unicast addresses (FE80::/64)
  - Multicast addresses (FF00::/8)
  - Unspecified address (0:0:0:0:0:0:0:0)

- Address with all F (F:F:F:F:F:F:F:F)

Caution After a DHCP server is created, you can update existing ranges, append new IP ranges, or delete existing ranges. However, it is a good practice to avoid deleting, shrinking, or expanding the existing IP ranges. For example, do not try to combine multiple smaller IP ranges to create a single large IP range. When you modify existing ranges after the DHCP service is running, it might cause the DHCP clients to lose network connection and result in a temporary traffic disruption.

Excluded Ranges (Only for DHCPv6)

Enter IPv6 addresses or a range of IPv6 addresses that you want to exclude for dynamic IP assignment to DHCPv6 clients.

In IPv6 networks, the DHCP ranges can be large. Sometimes, you might want to reserve certain IPv6 addresses, or multiple small ranges of IPv6 addresses from the large DHCP range for static binding. In such situations, you can specify excluded ranges.

Lease Time

Default value is 86400 seconds. Valid range of values is 60–4294967295.

Preferred Time (Only for DHCPv6)

Preferred time is the length of time that a valid IP address is preferred. When the preferred time expires, the IP address becomes deprecated. If no value is entered, preferred time is autocalculated as (lease time * 0.8).

Lease time must be > preferred time.

Valid range of values is 60–4294967295. Default is 69120 seconds.

DNS Servers

A maximum of two DNS servers are permitted. When not specified, no DNS is assigned to the DHCP client.

Domain names (Only for DHCPv6)

One or more domain names are supported.

DHCPv4 server configuration automatically fetches the domain name that you specified in the segment configuration.

SNTP Servers (Only for DHCPv6)

A maximum of two SNTP servers are permitted.

DHCPv6 server does not support NTP.

DHCP Options (Only for DHCPv4)

DHCP Options for IPv6 are not supported.

Each classless static route option in DHCP for IPv4 can have multiple routes with the same destination. Each route includes a destination subnet, subnet mask, next hop router. For information about classless static routes in DHCPv4, see RFC 3442 specifications. You can add a maximum of 127 classless static routes on a DHCPv4 server.

In addition to the Generic Option 121 (classless static route), NSX supports other Generic Options that are described in the following table. The Generic Options, which are not listed in this table are also accepted without any validation, but they do not take effect.

Table 5-2. Supported Generic Options

Code	Name	Value Type	Example Value
2	Time Offset	Integer - seconds offset from UTC Allowed values: -43200–43200 Maximum items: 1	28800
13	Boot File Size	Number of blocks. One block is 512 bytes. Integer values: 1–65535 Maximum items: 1	1385
19	Forward On/Off	IP forwarding Allowed values: [0, 1] 1 for on, 0 for off Maximum items: 1	0
26	MTU Interface	MTU for a given interface. Allowed values: 68–65535 Maximum items: 1	9600
28	Broadcast Address	IP address Maximum items: 1	10.10.10.255
35	ARP Timeout	Integer (seconds) Allowed values: 0–4294967295	360
40	NIS Domain	Text Maximum: 255 characters	vmware.com
41	NIS Servers	IP addresses in a preferred order Maximum items: 63	10.10.10.10
42	NTP Servers	IP addresses in a preferred order Maximum items: 63	10.10.10.11
44	NETBIOS Name Server	IP addresses in a preferred order Maximum items: 63	10.10.10.12
45	NETBIOS Dist Server	IP addresses in a preferred order Maximum items: 63	10.10.10.13

Table 5-2. Supported Generic Options (continued)

Code	Name	Value Type	Example Value
46	NETBIOS Node Type	Integer encoding of node type Allowed values: [1, 2, 4, 6] Maximum items: 4 1 = B-node - broadcast no WINS 2 = P-node - WINS only 4 = M-node - broadcast then WINS 8 = H-node - WINS then broadcast	2
47	NETBIOS Scope	String encoded according to RFC 1001/1002 Maximum: 255 characters	
58	Renewal Time	N/A - based on the lease time between 0–4294967295 Maximum items: 1	300
59	Rebinding Time	N/A - based on the lease time between 0–4294967295 Maximum items: 1	300
64	NIS+ Domain Name	Text (domain name)	vmware.com
65	NIS+ Server Address	IP addresses in a preferred order	10.10.10.10
66	Server Name	Text (server domain name) Maximum: 255 characters	10.10.10.253
67	Bootfile Name	Text (file name) Maximum: 255 characters	`/tftpboot/pxelinux/pxelinux.bin`
117	Name Service Search	Not natively supported with API Allowed values: [0, 6, 41, 44, 65] Maximum items: 5	6
119	Domain Search	One or more domain names. Each domain name must be enclosed in quotes and separated by commas.	vmware.com
150	TFTP server address	IP address	10.10.10.10
209	PXE Configuration File	Maximum: 255 characters	`configs/common`
210	PXE Path Prefix	Maximum: 255 characters	`/tftpboot/pxelinux/files/`
211	PXE Reboot Time	Allowed values: 0–4294967295	1800

DHCP Static Bindings

In a typical network environment, you have VMs that run services, such as FTP, email servers, application servers, and so on. You might not want the IP address of these VMs to change in your network. In this case, you can bind a static IP address to the MAC address of each VM (DHCP client). The static IP address must belong to the subnet (if any) that is configured on the segment, and it must not overlap with the DHCP IP ranges and DHCP server IP address.

DHCP static bindings are supported when you are configuring either a Segment DHCP server or a Gateway DHCP server on the segment. When a segment is using a DHCP Relay, you cannot configure static bindings.

On a DHCP for IPv4 server, static bindings are supported regardless of whether the segment uses a Segment DHCP or a Gateway DHCP configuration. On a DHCP for IPv6 server, static bindings are supported only when the segment uses a Segment DHCP configuration.

Static Binding Options Common to DHCPv4 and DHCPv6 Server

The following table describes the static binding options that are common to DHCP for IPv4 and DHCP for IPv6 servers.

Option	Description
Name	Enter a unique display name to identify each static binding. The name must be limited to 255 characters.
MAC Address	Required. Enter the MAC address of the DHCP client to which you want to bind a static IP address. The following validations apply to MAC address in static bindings: ■ MAC address must be unique in all the static bindings on a segment that uses a Segment DHCP server. ■ MAC address must be unique in all the static bindings across all the segments that are connected to the gateway and which use the Gateway DHCP server. For example, consider that you have 10 segments connected to a tier-1 gateway. You use a Gateway DHCP server for four segments (Segment1 to Segment4), and a Segment DHCP server for the remaining six segments (Segment5 to Segment10). Assume that you have a total of 20 static bindings across all the four segments (Segment1 to Segment4), which use the Gateway DHCP server. In addition, you have five static bindings in each of the other six segments (Segment5 to Segment10), which use a Segment DHCP server. In this example: ■ The MAC address in each of the 20 static bindings must be unique across all the segments (Segment1 to Segment4) that use the Gateway DHCP server. ■ The MAC address in the five static bindings must be unique on each segment (Segment5 to Segment10) that use a Segment DHCP server.

Option	Description
IP Address	■ Required for IPv4 static binding. Enter a single IPv4 address to bind to the MAC address of the client. ■ Optional for IPv6 static binding. Enter a single Global Unicast IPv6 address to bind to the MAC address of the client. When no IPv6 address is specified for static binding, Stateless Address Autoconfiguration (SLAAC) is used to auto-assign an IPv6 address to the DHCPv6 clients. Also, you can use Stateless DHCP to assign other DHCP configuration options, such as DNS, domain names, and so on, to the DHCPv6 clients. For more information about Stateless DHCP for IPv6, read the RFC 3376 specifications. The following types of IPv6 addresses are not permitted in IPv6 static binding: ■ Link Local Unicast addresses (FE80::/64) ■ Multicast IPv6 addresses (FF00::/8) ■ Unspecified address (0:0:0:0:0:0:0:0) ■ Address with all F (F:F:F:F:F:F:F:F) The static IP address must belong to the subnet (if any) that is configured on the segment, and it must be outside the DHCP ranges that you have configured on the segment.
Lease Time	Optional. Enter the amount of time in seconds for which the IP address is bound to the DHCP client. When the lease time expires, the IP address becomes invalid and the DHCP server can assign the address to other DHCP clients on the segment. Valid range of values is 60–4294967295. Default is 86400.
Description	Optional. Enter a description for the static binding.
Tags	Optional. Add tags to label static bindings so that you can quickly search or filter bindings, troubleshoot and trace binding-related issues, or do other tasks. For more information about adding tags and use cases for tagging objects, see Tags.

Static Binding Options (Only in DHCPv4 Server)

The following table describes the static binding options that are available only in a DHCP for IPv4 server.

DHCP For IPv4 Option	Description
Gateway Address	Enter the default gateway IP address that the DHCP for IPv4 server must provide to the DHCP client.
Host Name	Enter the host name of the DHCP for IPv4 client so that the DHCPv4 server can always bind the client (host) with the same IPv4 address each time. The host name must be limited to 63 characters. The following validations apply to host name in static bindings: ■ Host name must be unique in all the static bindings on a segment that uses a Segment DHCP server. ■ Host name must be unique in all the static bindings across all the segments that are connected to the gateway and which use the Gateway DHCP server. For example, consider that you have 10 segments connected to a tier-1 gateway. You use a Gateway DHCP server for four segments (Segment1 to Segment4), and a Segment DHCP server for the remaining six segments (Segment5 to Segment10). Assume that you have a total of 20 static bindings across all the four segments (Segment1 to Segment4), which use the Gateway DHCP server. In addition, you have five static bindings in each of the other six segments (Segment5 to Segment10), which use a Segment DHCP server. In this example: ■ The host name in each of the 20 static bindings must be unique across all the segments (Segment1 to Segment4) that use the Gateway DHCP server. ■ The host name in the five static bindings must be unique on each segment (Segment5 to Segment10) that use a Segment DHCP server.
DHCP Options	Optional. Click **Set** to configure DHCP for IPv4 Classless Static Routes and other Generic Options.

Some additional notes for DHCPv4 static binding:

- IPv4 static bindings automatically inherit the domain name that you configured on the segment.
- To specify DNS servers in the static binding configuration, add the **Generic Option** (Code 6 - DNS Servers).
- To synchronize the system time on DHCPv4 clients with DHCPv4 servers, use NTP. DHCP for IPv4 server does not support SNTP.
- If DHCP options are not specified in the static bindings, the DHCP options from the DHCPv4 server on the segment are automatically inherited in the static bindings. However, if you have explicitly added one or more DHCP options in the static bindings, these DHCP options are not autoinherited from the DHCPv4 server on the segment.

Static Binding Options (Only in DHCPv6 Server)

The following table describes the static binding options that are available only in a DHCP for IPv6 server.

DHCP for IPv6 Option	Description
DNS Servers	Optional. Enter a maximum of two domain name servers to use for the name resolution. When not specified, no DNS is assigned to the DHCP client.
SNTP Servers	Optional. Enter a maximum of two Simple Network Time Protocol (SNTP) servers. The clients use these SNTP servers to synchronize their system time to that of the standard time servers.
Preferred Time	Optional. Enter the length of time that a valid IP address is preferred. When the preferred time expires, the IP address becomes deprecated. If no value is entered, preferred time is auto-calculated as (lease time * 0.8). Lease time must be > preferred time. Valid range of values is 60–4294967295. Default is 69120.
Domain Names	Optional. Enter the domain name to provide to the DHCPv6 clients. Multiple domain names are supported in an IPv6 static binding. When not specified, no domain name is assigned to the DHCP clients.

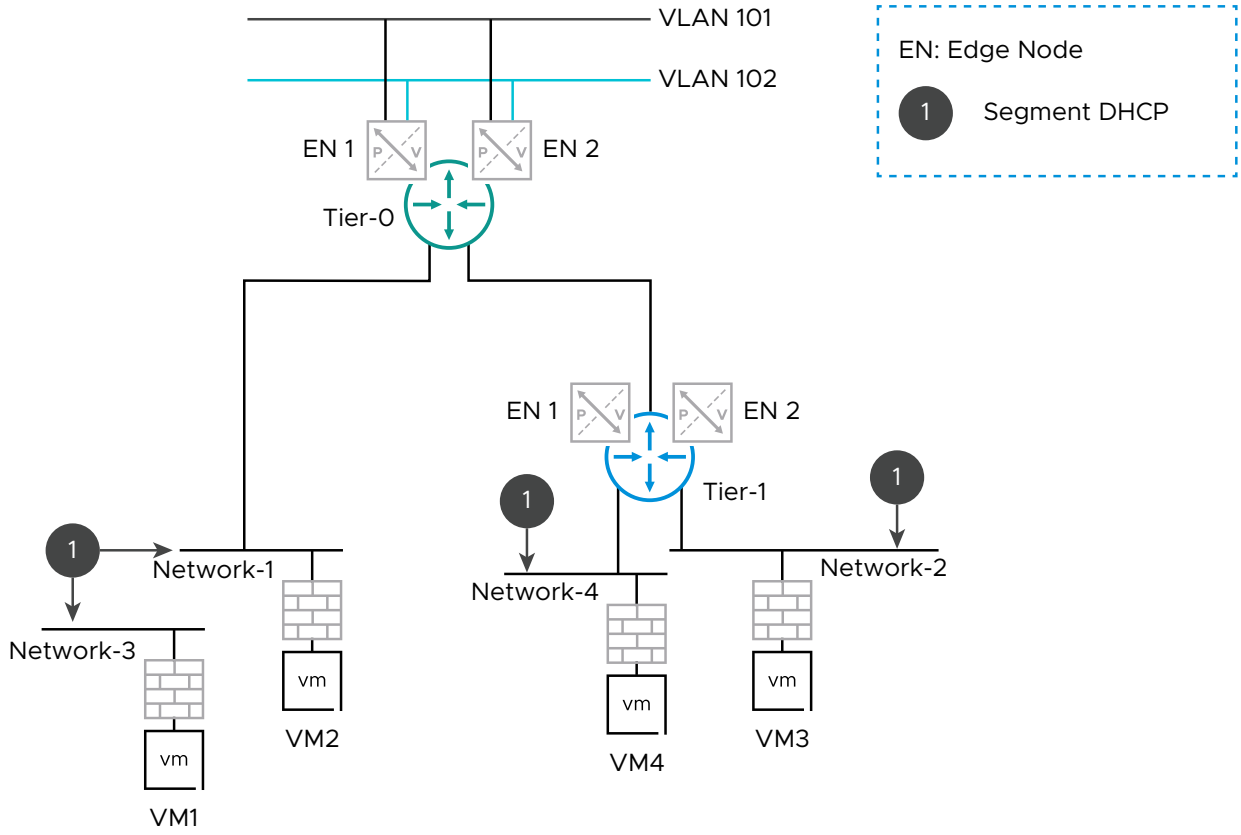
Some additional notes for DHCPv6 static binding:

- Gateway IP address configuration is unavailable in IPv6 static bindings. IPv6 client learns about its first-hop router from the ICMPv6 router advertisement (RA) message.
- NTP is not supported in DHCPv6 static bindings.

Configure Segment DHCP Server on a Segment

A Segment DHCP server provides a dynamic IP assignment service only to the VMs that are attached to the segment. NSX supports Segment DHCP server configuration on the downlink interface and the service interface. You can configure a Segment DHCPv4 server, or a Segment DHCPv6 server, or both, on the segment.

The following figure shows a sample network topology that has a Segment DHCP server configured on four networks.



In this network topology, a Segment DHCP server is configured on the following networks:

- Network-2 is connected to the downlink interface of tier-1 gateway.
- Network-1 is connected to the service interface of tier-0 gateway.
- Network-4 is connected to the service interface of tier-1 gateway.
- Network-3 is an isolated segment, which is not connected to any gateway.

Procedure

1. From your browser, log in with **admin** privileges to an NSX Manager at https://*nsx-manager-ip-address*.
2. Select **Networking** > **Segments**.
3. Find the segment where you want to configure the DHCP service. Next to the segment name, click , and then click **Edit**.
4. Click **Set DHCP Config**.
5. From the **DHCP Type** drop-down menu, select **Segment DHCP Server**.

6 (Required) From the **DHCP Profile** drop-down menu, select a DHCP server profile. If no profile is available in the drop-down menu, click  and then click **Create New** to add a DHCP server profile. After the profile is created, it is automatically attached to the segment.

   For more information about creating a DHCP server profile, see Add a DHCP Server Profile.

7 Click the **DHCP Config** toggle button to enable DHCP configuration on the segment.

   If you are configuring a Segment DHCPv4 server and a Segment DHCPv6 server, ensure that you enable the **DHCP Config** toggle button in both the **IPv4 Server** and **IPv6 Server** tabs.

8 Specify the following DHCP configuration settings:

   - DHCP Server Address
   - DHCP Ranges
   - Optional: Excluded Ranges (only for DHCPv6)
   - Optional: Lease Time
   - Optional: Preferred Time (only for DHCPv6)
   - Optional: Domain Names (only for DHCPv6)
   - Optional: DNS Servers
   - Optional: SNTP Servers (only for DHCPv6)
   - Optional: DHCP Options (only for DHCPv4)
   - Optional: Static Bindings

   For a detailed information about these DHCP configuration settings, see the reference documentation at DHCP Configuration Settings: Reference.

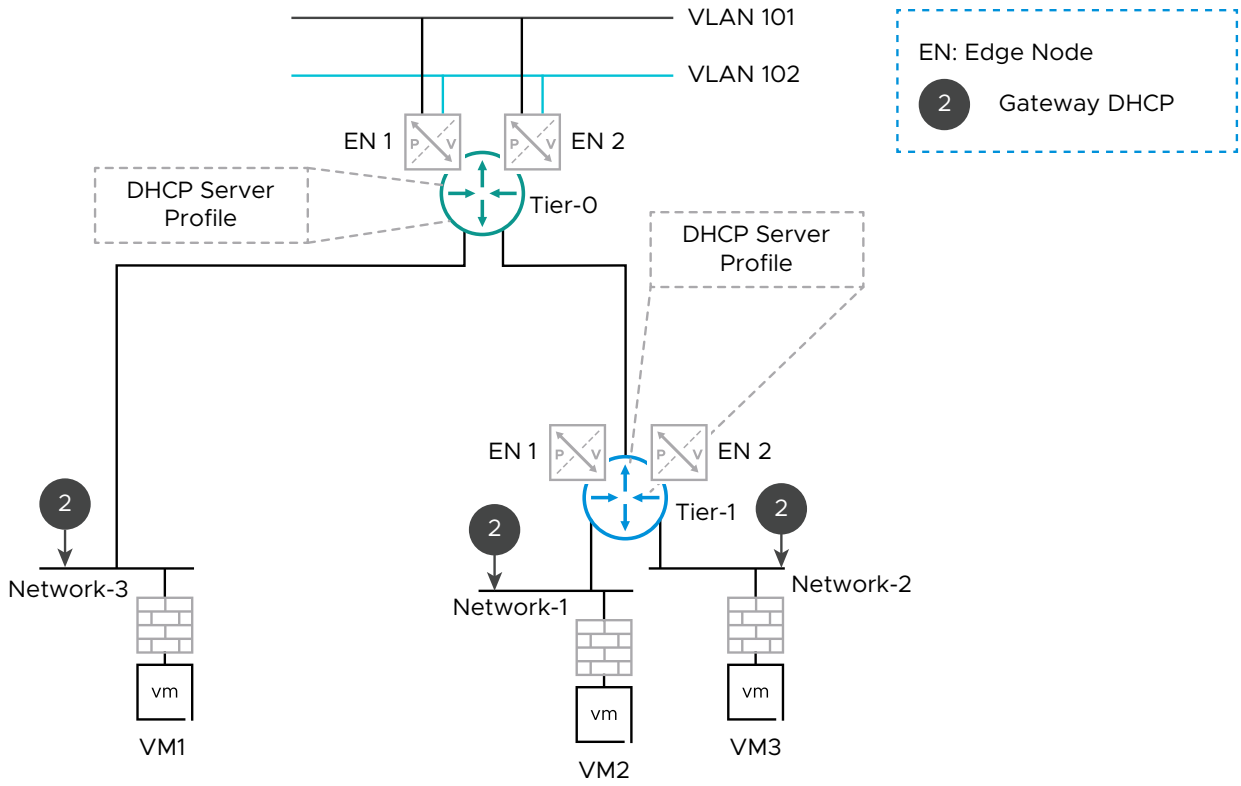
9 Click **Apply**.

Configure Gateway DHCP Server on a Segment

A Gateway DHCP server is attached to a tier-0 or tier-1 gateway, and it provides DHCP service to the networks (overlay segments), which are directly connected to the gateway and configured to use a Gateway DHCP server.

In this case, the DHCP server that is created on the tier-0 or tier-1 gateway will have an internal relay so that the connected segments can forward traffic to the DHCP servers, which you specified in the DHCP server profile.

The following figure shows a sample network topology with Gateway DHCP servers that are servicing networks, which are directly connected to the tier-0 and tier-1 gateway.



In this network topology, a Gateway DHCP server is configured on the following networks:

- Network-1 is connected to the service interface of the tier-1 gateway.
- Network-2 is connected to the downlink interface of the tier-1 gateway.
- Network-3 is connected to the downlink or service interface of the tier-0 gateway.

Prerequisites

Ensure that you have specified the Gateway IP address of the IPv4 subnet in the segments that are directly connected to the tier-0 or tier-1 gateway.

Procedure

1. From your browser, log in with **admin** privileges to an NSX Manager at https://*nsx-manager-ip-address*.
2. Select **Networking** > **Segments**.
3. Find the segment where you want to configure the DHCP service. Next to the segment name, click , and then click **Edit**.
4. Click **Set DHCP Config**.
5. In the **DHCP Type** drop-down menu, select **Gateway DHCP Server**.

6 (Required) Ensure that a DHCP server profile is attached to the gateway.

If a profile is set on the gateway, the name of the profile is displayed in a read-only mode.

If a DHCP server profile is not set on the gateway, do these steps:

a Click the information icon next to **DHCP Profile**, and then click the gateway name to navigate to the gateway page.

b Next to the gateway name, click  , and then click **Edit**.

c Next to **DHCP Config**, click **Set**.

The **Set DHCP Configuration** window opens.

d In the **Type** drop-down menu, select **DHCP Server**.

e Select a DHCP server profile to attach to this gateway and click **Save**.

f Close the edit mode on the gateway and return to the edit mode on the Segments page.

g Click **Set DHCP Config** and continue the remaining steps in this procedure.

7 Click the **DHCP Config** toggle button to enable DHCP configuration on the segment.

Note You can configure only Gateway DHCPv4 server on a segment. Gateway DHCPv6 server is not supported.

8 Observe that the **DHCP Server Address** is fetched automatically from the DHCP profile and displayed on the **IPv4 Server > Settings** page.

9 Specify the following DHCP configuration settings:

- DHCP Ranges
- Optional: Lease Time
- Optional: DNS Servers
- Optional: DHCP Options
- Optional: Static Bindings

For a detailed information about these DHCP configuration settings, see the reference documentation at DHCP Configuration Settings: Reference.

10 Click **Apply**.

Configure DHCP Relay on a Segment

In a DHCP Relay configuration, the DHCP messages are forwarded to the external DHCP servers. The external DHCP servers can be in any subnet, outside the SDDC, or in the physical network.

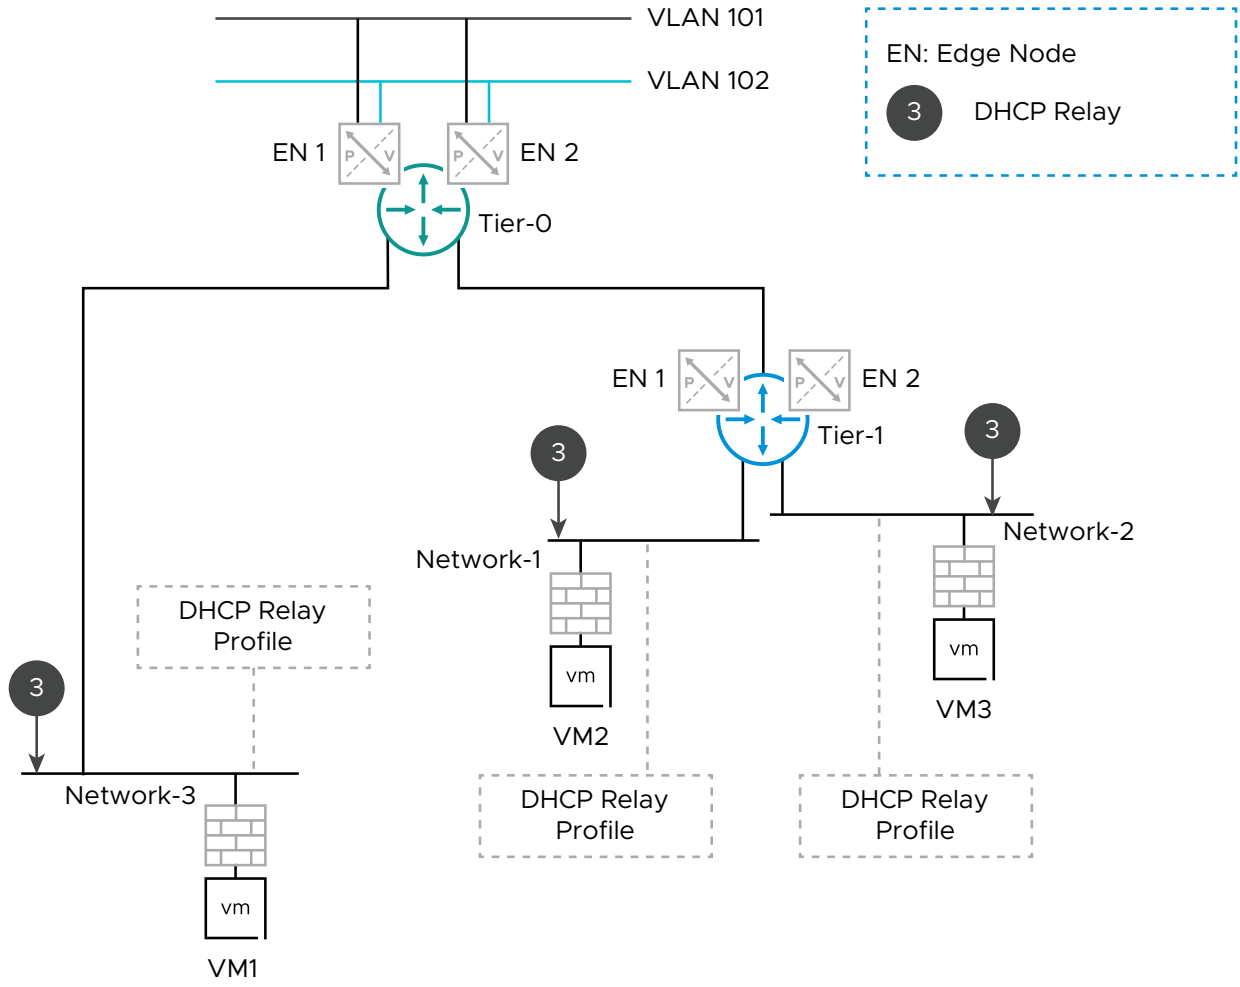
DHCP Relay configuration is supported in the following scenarios:

- When an overlay segment is connected to the downlink interface of a tier-0 or tier-1 gateway. In this case, the DHCP messages can be relayed either to DHCPv4 servers or DHCPv6 servers. Step 2 in the **Procedure** section of this topic explains the workflow for this scenario.

- When an overlay or VLAN segment is connected to the service interface of a tier-0 or tier-1 gateway that is configured in an active-standby mode. In this case, the DHCP messages are relayed only to DHCPv4 servers. Step 3 in the **Procedure** section of this topic explains the workflow for this scenario.

Note When you use a DHCP Relay on a segment, you cannot configure DHCP settings, DHCP options, and static bindings on the segment.

The following figure shows a sample network topology that has a DHCP Relay configured on three networks.



In this network topology, a DHCP Relay is configured on the following networks:

- Network-1 is connected to the service interface of the tier-1 gateway.
- Network-2 is connected to the downlink interface of the tier-1 gateway.
- Network-3 is connected either to the downlink or service interface of the tier-0 gateway.

Procedure

1 From your browser, log in with **admin** privileges to an NSX Manager at https://*nsx-manager-ip-address*.

2 Configure a DHCP Relay on an overlay segment that is connected to the downlink interface of a tier-0 or tier-1 gateway.

   a Navigate to **Networking > Segments**.

   b Find the overlay segment where you want to configure the DHCP Relay. Next to the segment name, click  , and then click **Edit**.

   c Click **Set DHCP Config**.

   d From the **DHCP Type** drop-down menu, select **DHCP Relay**.

   e From the **DHCP Profile** drop-down menu, select a DHCP relay profile. If no profile is available in the drop-down menu, click  and then click **Create New** to add a DHCP relay profile. After the profile is created, it is automatically attached to the segment.

   f Click **Apply**.

3 Configure a DHCP Relay on a segment that is connected to the service interface of a tier-0 or tier-1 gateway, which is configured in an active-standby mode.

   a Navigate to **Networking > Segments**.

   b Add a segment in either a VLAN or an overlay transport zone. Do not connect this segment to any gateway. Also, do not set any DHCP configuration on this segment, such as DHCP server address, DHCP ranges, and static bindings.

   For example, assume that you have added a segment in the VLAN transport zone with name as **`My-VLAN-Segment`**.

   c Navigate to **Networking > Tier-0 Gateways** or **Networking > Tier-1 Gateways**.

   d Find the gateway where you want to connect this VLAN segment to the service interface. Click  , and then click **Edit**.

   e Expand the **Interfaces** section and click the link to open the **Set Interfaces** page.

   f Click **Add Interface**.

   g In the **Name** text box, enter a name for this interface.

   For example, specify the name as **`Connect-to-VLAN`**.

   h From the **Type** drop-down menu, select **Service**.

   i Enter the **IP Address/Mask** in a CIDR format.

   For example, enter **`172.16.10.1/24`**.

j From the **Connect To (Segment)** drop-down menu, select the **My-VLAN-Segment**, which you created earlier.

k From the **DHCP Profile** drop-down menu, select a DHCP relay profile.

l Click **Close**.

Attach a DHCP Profile to a Tier-0 or Tier-1 Gateway

To use Gateway DHCP server for a dynamic IP assignment, you must attach a DHCP server profile to a tier-0 or tier-1 gateway.

You can attach a DHCP profile to a gateway only when the segments connected to that gateway do not have a Segment DHCP server or a DHCP relay configured on them. If a Segment DHCP server or DHCP relay exists on the segment, the UI throws an error when you try to attach a DHCP profile to the gateway. You must disconnect the segments from the gateway, and then attach a DHCP profile to the gateway.

Prerequisites

A DHCP server profile is added in the network.

Procedure

1 From your browser, log in with **admin** privileges to an NSX Manager at `https://nsx-manager-ip-address`.

2 Go to **Networking > Tier-0 Gateways** or **Networking > Tier-1 Gateways**.

3 Edit the appropriate gateway.

4 Next to **DHCP Config**, click **Set**.

The **Set DHCP Configuration** window opens.

5 In the **Type** drop-down menu, select **DHCP Server**.

6 Select a DHCP server profile to attach to this gateway.

7 Click **Save**.

What to do next

Navigate to **Networking > Segments**. On each segment that is connected to this gateway, configure the DHCP settings, DHCP options, and static bindings.

For more information, see Configure DHCP Service.

View Gateway DHCP Statistics

After a Gateway DHCP server is in use, you can view the DHCP server statistics on the directly connected tier-0 or tier-1 gateway.

Prerequisites

- DHCP server profile is attached to the tier-0 or tier-1 gateway.
- Segments that are connected to the gateway are configured to use a Gateway DHCP server.
- DHCP settings are configured on the segments that are directly connected to the gateway.
- The server runtime status is up and the Gateway DHCP server is in use.

Procedure

1 From your browser, log in to an NSX Manager at https://*nsx-manager-ip-address*.

2 Navigate to **Networking > Tier-0 Gateways** or **Networking > Tier-1 Gateways**.

3 Find the gateway whose Gateway DHCP server statistics you want to view.

4 Expand the gateway configuration settings, and then click the link next to **DHCP**.

5 In the pop-up window, click **View Statistics**.

Gateway DHCPv4 server statistics are displayed.

In the DHCP Server Packets section, a pie chart displays the breakup of the DHCP packet counts. These packet counts represent the count of the various DHCP message types. The number at the center of the pie chart represents the sum of all DHCP packet counts.

The DHCP Pool Statistics section displays the pool usage statistics of segments that are directly connected to the gateway and using Gateway DHCP server. For example, this section shows statistics, such as the size of the DHCP pool, the number of IP addresses used from the pool, and the allocation percentage.

Note If you have configured a Segment DHCP server on a gateway-connected segment, the statistics of the Segment DHCP server are not displayed on the **DHCP Statistics** page of the gateway. DHCP statistics are displayed only for those segments that are configured to use a Gateway DHCP server.

For example, assume that you have four segments connected to the downlink interfaces of a tier-1 gateway. Segments 1 and 2 are using a Segment DHCP server, whereas segments 3 and 4 are using the Gateway DHCP server. In this case, the **DHCP Statistics** page displays the Gateway DHCP server statistics only for segments 3 and 4.

To view segment DHCP statistics, you must navigate to the **Segments** page. For more information, see View Segment DHCP Statistics.

6 (Optional) To reset DHCP packet counts, click **Reset Packet Counter**.

The DHCP packet counts that are displayed next to the pie chart are reset. The DHCP pool statistics are not impacted.

View Segment DHCP Statistics

After a Segment DHCP server is in use, you can view the DHCP server statistics on the **Segments** page.

Prerequisites

- DHCP settings are configured on the segment.
- The server runtime status is up and the Segment DHCP server is in use.

Procedure

1. From your browser, log in to an NSX Manager at `https://nsx-manager-ip-address`.
2. Select **Networking > Segments**.
3. Find the segment whose DHCP statistics you want to view.
4. Expand the segment configuration settings, and then click **View Statistics**.

   The **Segment Statistics** page opens.
5. Click the **DHCP Statistics** tab.

   In the DHCP Server Packets section, a pie chart displays the breakup of the DHCP packet counts. The packet counts represent the count of the various DHCP message types. The number at the center of the pie chart represents the sum of all DHCP packet counts.

   The DHCP Pool Statistics section displays the pool usage statistics. For example, this section shows statistics, such as the size of the DHCP pool, the number of IP addresses used from the pool, and the allocation percentage.

   Important If you have configured both DHCPv4 and DHCPv6 servers on a segment, the **DHCP Statistics** page will display only the DHCPv4 packet counts and the DHCPv4 pool usage statistics. DHCPv6 packet counts and DHCPv6 pool usage statistics are currently not supported.
6. (Optional) To reset DHCP packet counts, click **Reset Packet Counter**.

   The DHCP packet counts that are displayed next to the pie chart are reset. The DHCP pool statistics are not impacted.

Scenarios: Selection of Edge Cluster for DHCP Service

DHCP server runs as a service (service router) in the edge nodes of an NSX Edge cluster.

Isolated segments that are not connected to a gateway can use only a Segment DHCP server. Segments that are connected to a gateway on the downlink interface can use either a Segment DHCP server, DHCP Relay, or Gateway DHCP server.

Regardless of whether a segment uses a Segment DHCP server or a Gateway DHCP server, DHCP server always runs as a service router in the edge transport nodes of an edge cluster. If the segment uses a Segment DHCP server, the DHCP service is created in the edge cluster that you specified in the DHCP profile. However, if the segment uses a Gateway DHCP server, the edge cluster in which the DHCP service is created depends on the combination of the following factors:

- Is an edge cluster specified in the gateway?
- Is an edge cluster specified in the DHCP profile of the gateway?
- Is the edge cluster in the gateway and in the DHCP profile same or different?
- Is the tier-1 routed segment connected to a tier-0 gateway?

The following scenarios explain how the edge cluster is selected for creating the DHCP service.

Scenario 1: Isolated Segment Uses Segment DHCP Server

Scenario Description:

- An edge cluster (Cluster1) is created with four edge nodes: N1, N2, N3, N4.
- A segment with None connectivity is added in the overlay transport zone.
- Segment uses a Segment DHCP server, by default.

The DHCP server profile configuration is as follows:

- Profile Type: **`DHCP Server`**
- Edge Cluster: **`Cluster1`**
- Edges: **`Autoallocated`**

In this scenario, any two edge nodes from Cluster1 are autoallocated to create the DHCP service, and DHCP high availability (HA) is automatically configured. One of the edge nodes in Cluster1 runs in active mode and the other edge runs in passive mode.

Note

- If you manually allocate the edge nodes in the DHCP profile, the edge node that is added first becomes the active edge. The second edge node takes the passive role.
- If you select only one edge node in the DHCP profile, DHCP HA is not configured.

Scenario 2: Tier-1 Routed Segment Uses Gateway DHCP and Different Edge Clusters in Gateway and DHCP Profile

Consider that you have two edge clusters in your network (Cluster1 and Cluster2). Both clusters have four edge nodes each:

- Cluster1 edge nodes: N1, N2, N3, N4
- Cluster2 edge nodes: N5, N6, N7, N8

Scenario Description:

- Segment is connected to a tier-1 gateway.
- Tier-1 gateway is not connected to a tier-0 gateway.
- DHCP server profile in the tier-1 gateway uses Cluster1.
- Tier-1 gateway uses Cluster2.
- Segment is configured to use the Gateway DHCP server.

The DHCP server profile in the tier-1 gateway has the following configuration:

- Profile Type: **`DHCP Server`**
- Edge Cluster: **`Cluster1`**
- Edges: **`N1,N2`** (allocated manually in the given sequence)

The tier-1 gateway configuration is as follows:

- Edge Cluster: **`Cluster2`**
- Edges: **`N5,N6`** (allocated manually in the given sequence)

In this scenario, DHCP service runs on the edge nodes of Cluster2. As Cluster2 contains multiple edge nodes, DHCP HA is autoconfigured. However, the manually allocated edges N5 and N6 on the gateway are ignored for DHCP HA. Any two nodes from Cluster2 are randomly autoallocated for DHCP HA.

This scenario also applies when the segment is directly connected to a tier-0 gateway, and there is no tier-1 gateway in your network topology.

Caution You can change the edge cluster on the Gateway DHCP server after the DHCP server is created. However, this action causes all the existing DHCP leases that are assigned to the DHCP clients to be lost.

To summarize, the main points of this scenario are as follows:

- When you use a Gateway DHCP server and set different edge clusters in the gateway DHCP profile and tier-1 gateway, then DHCP service is always created in the edge cluster of the gateway.
- The edge nodes are randomly allocated from the edge cluster of the tier-1 gateway for DHCP HA configuration.
- If no edge cluster is specified in the tier-1 gateway, the edge cluster in the DHCP profile of the tier-1 gateway (Cluster1) is used to create the DHCP service.

Scenario 3: Tier-1 Routed Segment Uses Segment DHCP Server and Different Edge Clusters in Gateway and DHCP Profile

In this scenario, a segment is connected to a tier-1 gateway, but you use a Segment DHCP server on the segment. Consider that you have three edge clusters in your network (Cluster1, Cluster2, Cluster 3). Each cluster has two edges nodes each.

- Cluster1 edge nodes: N1, N2
- Cluster2 edge nodes: N3, N4
- Cluster3 edge nodes: N5, N6

Scenario Description:

- Segment is connected to a tier-1 gateway.
- Tier-1 gateway is connected to a tier-0 gateway (optional).
- DHCP profile on the gateway uses Cluster1.
- Gateway uses Cluster2.
- Segment is configured to use Segment DHCP server.
- DHCP server profile on the segment uses Cluster3.

The DHCP server profile on the gateway is as follows:

- Profile Name: **`ProfileX`**
- Profile Type: **`DHCP Server`**
- Edge Cluster: **`Cluster1`**
- Edges: **`N1,N2`** (allocated manually in the given sequence)

The tier-1 gateway configuration is as follows:

- Edge Cluster: **`Cluster2`**
- Edges: **`N3,N4`** (allocated manually in the given sequence)

The profile on the Segment DHCP server is as follows:

- Profile Name: **`ProfileY`**
- Profile Type: **`DHCP Server`**
- Edge Cluster: **`Cluster3`**
- Edges: **`N5,N6`** (allocated manually in the given sequence)

In this scenario, because the segment is configured to use a Segment DHCP server, the edge cluster (Cluster2) in the connected tier-1 gateway is ignored to create the DHCP service. DHCP service runs in the edge nodes of Cluster3 (N5, N6). DHCP HA is also configured. N5 becomes the active edge node and N6 becomes the standby edge.

If edge nodes are not manually allocated from Cluster3, any two nodes from this cluster are autoallocated for creating the DHCP service and configuring DHCP HA. One of the edge nodes becomes an active edge and the other node becomes the standby edge. If only one edge node is allocated manually from Cluster3, DHCP HA is not configured.

This scenario also applies when the segment is directly connected to a tier-0 gateway, and there is no tier-1 gateway in your network topology.

Scenario 4: Tier-1 Routed Segment Uses Gateway DHCP and Same Edge Clusters in Gateway and DHCP Profile

Consider that you have a single edge cluster (Cluster1) in your network with four edge nodes: N1, N2, N3, N4.

Scenario Description:

- Segment is connected to a tier-1 gateway.
- Tier-1 gateway is connected to a tier-0 gateway (optional)
- Gateway and DHCP profile on the gateway use the same edge cluster (Cluster1).
- Segment is configured to use Gateway DHCP server.

The DHCP server profile on the gateway is as follows:

- Profile Type: **`DHCP Server`**
- Edge Cluster: **`Cluster1`**
- Edges: **`N1,N2`** (allocated manually in the given sequence)

The tier-1 gateway configuration is as follows:

- Edge Cluster: **`Cluster1`**
- Edges: **`N3,N4`** (allocated manually in the given sequence)

In this scenario, as the gateway DHCP profile and gateway use a similar edge cluster (Cluster1), DHCP service is created in the edge nodes N1 and N2 of the gateway DHCP profile. The edge nodes N3 and N4 that you specified in the connected tier-1 gateway are ignored for creating the DHCP service.

If edge nodes are not manually set in the DHCP profile, any two nodes from Cluster1 are autoallocated for creating the DHCP service and configuring DHCP HA. One of the edge nodes becomes an active edge and the other edge becomes the standby edge.

To summarize, the main points of this scenario are as follows:

- When you use a Gateway DHCP server and specify similar edge clusters in the DHCP profile and connected gateway, then DHCP service is created in the edge nodes of the DHCP profile.
- The edges nodes that you manually specified in the connected gateway are ignored.

Scenario 5: Tier-1 Routed Segment is Connected to Tier-0 Gateway and No Edge Cluster is Set in Tier-1 Gateway

In this scenario, a segment is connected to a tier-1 gateway, and the tier-1 gateway is connected to a tier-0 gateway. Consider that you have three edge clusters in your network (Cluster1, Cluster2, Cluster 3). Each cluster has two edges nodes each.

- Cluster1 edge nodes: N1, N2
- Cluster2 edge nodes: N3, N4
- Cluster3 edge nodes: N5, N6

Scenario Description:

- Segment is directly connected to a tier-1 gateway.
- Tier-1 gateway is connected to a tier-0 gateway.
- DHCP server profile is specified on both tier-1 and tier-0 gateways.
- DHCP profile on tier-1 gateway uses Cluster1.
- DHCP profile on tier-0 gateway uses Cluster2.
- No edge cluster is selected in tier-1 gateway.
- Tier-0 gateway uses Cluster3.
- Segment is configured to use a Gateway DHCP server.

In this scenario, because the tier-1 gateway has no edge cluster specified, NSX falls back on the edge cluster of the connected tier-0 gateway. DHCP service is created in the edge cluster of tier-0 gateway (Cluster3). Any two edge nodes from this edge cluster are autoallocated for creating the DHCP service and configuring DHCP HA.

To summarize, the main points of this scenario are as follows:

- When a tier-1 gateway has no edge cluster specified, NSX falls back on the edge cluster of the connected tier-0 gateway to create the DHCP service.
- If no edge cluster is detected in the tier-0 gateway, DHCP service is created in the edge cluster of the tier-1 gateway DHCP profile.

Scenarios: Impact of Changing Segment Connectivity on DHCP

After you save a segment with DHCP configuration, you must be careful about changing the connectivity of the segment.

Segment connectivity changes are allowed only when the segments and gateways belong to the same transport zone.

The following scenarios explain the segment connectivity changes that are allowed or disallowed, and whether DHCP is impacted in each of these scenarios.

Scenario 1: Move a Routed Segment with Gateway DHCP Server to a Different Gateway

Consider that you have added a segment and connected it either to a tier-0 or tier-1 gateway. You configured Gateway DHCP server on this segment, saved the segment, and connected workloads to this segment. DHCP service is now used by the workloads on this segment.

Later, you decide to change the connectivity of this segment to another tier-0 or tier-1 gateway, which is in the same transport zone. This change is allowed. However, when you save the segment, an information message alerts you that changing the gateway connectivity impacts the existing DHCP leases, which are assigned to the workloads.

Scenario 2: Move a Routed Segment with Segment DHCP Server or Relay to a Different Gateway

Consider that you have added a segment and connected it either to a tier-0 or tier-1 gateway. You configured Segment DHCP server or DHCP Relay on this segment, saved the segment, and connected workloads to this segment. DHCP service is now used by the workloads on this segment.

Later, you decide to change the connectivity of this segment to another tier-0 or tier-1 gateway, which is in the same transport zone. This change is allowed. As the DHCP server is local to the segment, the DHCP configuration settings, including ranges, static bindings, and DHCP options are retained on the segment. The DHCP leases of the workloads are retained and there is no loss of network connectivity.

After the segment is moved to a new gateway, you can continue to update the DHCP configuration settings and other segment properties. You can change the DHCP type and DHCP profile of a routed segment after moving the segment to a different gateway.

Scenario 3: Move a Standalone Segment with Segment DHCP Server to a Tier-0 or Tier-1 Gateway

Consider that you have added a segment with None connectivity in your network. You have configured Segment DHCP server on this segment, saved the segment, and connected workloads to this segment. DHCP service is now used by the workloads on this segment.

Later, you decide to connect this segment either to a tier-0 or tier-1 gateway, which is in the same transport zone. This change is allowed. As a Segment DHCP server existed on the segment, the DHCP configuration settings, including ranges, static bindings, and DHCP options are retained on the segment. The DHCP leases of the workloads are retained and there is no loss of network connectivity.

After the segment is connected to the gateway, you can continue to update the DHCP configuration settings, and other segment properties. However, you cannot select a different DHCP type and the DHCP profile in the segment. For example, you cannot change the DHCP type from a Segment DHCP server to a Gateway DHCP server or a DHCP Relay. In addition, you cannot change the DHCP server profile in the segment. But you can edit the properties of the DHCP profile, if needed.

Scenario 4: Move a Standalone Segment Without DHCP Configuration to a Tier-0 or Tier-1 Gateway

Consider that you have added a segment with None connectivity in your network. You have not configured DHCP on this segment, saved the segment, and connected workloads to this segment.

Later, you decide to connect this segment either to a tier-0 or tier-1 gateway, which is in the same transport zone. This change is allowed. As no DHCP configuration existed on the segment, the segment automatically uses the Gateway DHCP server after it is connected to the gateway. The DHCP profile attached to this gateway gets autoselected in the segment.

Now, you can specify the DHCP configuration settings, including ranges, static bindings, and DHCP options on the segment. You can also edit the other segment properties, if necessary. However, you cannot change the DHCP type from a Gateway DHCP server to a Segment DHCP server or a DHCP Relay.

Remember, you can configure only a Gateway DHCPv4 server on the segment. Gateway DHCPv6 server is not supported.

Scenario 5: Move a Segment with Tier-0 or Tier-1 Connectivity to None Connectivity

Consider that you have added a segment to a tier-0 or tier-1 gateway in your network. You have configured Gateway DHCP server or DHCP Relay on this segment, saved the segment, and connected workloads to this segment. DHCP service is now used by the workloads on this segment.

Later, you decide to change the connectivity of this segment to None. This change is not allowed.

In this scenario, the following workaround can help:

1 Temporarily disconnect the existing segment from the gateway or delete the segment.
   a In NSX Manager, navigate to **Networking > Segments**.
   b Click the vertical ellipses next to the segment, and then click **Edit**.
   c Turn off the **Gateway Connectivity** option to disconnect the segment temporarily from the gateway.
2 Add a new segment and do not connect it to any gateway.
3 Configure a Segment DHCP server on this standalone segment, if needed.

Host Switches

6

A host switch managed object is a virtual network switch that provides networking services to the various hosts in the network. It is instantiated on every host that participates in NSX networking

The following host switches are supported in NSX:

- NSX Virtual Distributed Switch: NSX introduces a host switch that normalizes connectivity among various compute domains, including multiple VMware VMware vCenter instances, containers, and other off premises or cloud implementations.
- NSX Virtual Distributed Switch can be configured based on the performance required in your environment:
  - Standard: Configured for regular workloads, where normal traffic throughput is expected on the workloads.
  - Enhanced: Configured for telecom workloads, where high traffic throughput is expected on the workloads.
- vSphere Distributed Virtual Switch: Provides centralized management and monitoring of the networking configuration of all hosts that are associated with the switch in a VMware vCenter environment.

This chapter includes the following topics:

- Managing NSX on a vSphere Distributed Switch
- Enhanced Datapath
- Configure VTEP HA
- VTEP Maintenance Mode

Managing NSX on a vSphere Distributed Switch

Configure and run NSX on a vSphere Distributed Switch (VDS).

In NSX 4.0, you can only use a VDS switch to prepare ESXi host nodes as transport nodes. N-VDS switch is not supported. Configure the NSX Distributed Firewall functionality on VDS in data centers and workloads where segmentation, visibility, or advanced security capabilities are desired. This ensures distributed firewall capabilities work on a VM managed by whether it is managed by vCenter Server.

However, to prepare an NSX Edge VM as a transport node, you can only use a N-VDS switch. You can connect a NSX Edge VM to any of the supported host switches (VSS or VDS) depending on the topology in your network.

After you prepare a cluster of transport node hosts with VDS as the host switch, you can do the following:

- Manage NSX transport nodes on a VDS switch.
- Realize a segment created in NSX as an NSX Distributed Virtual port group in vCenter Server.
- Migrate VMs between VDS port groups.

Configuring a vSphere Distributed Switch

When a transport node is configured on a VDS host switch, some network parameters can only be configured in VMware vCenter.

The following requirements must be met to install NSX on a VDS host switch:

- VMware vCenter 7.0 or a later version
- ESXi 7.0 or a later version

The created VDS switch can be configured to centrally manage networking for NSX hosts.

Configuring a VDS switch for NSX networking requires objects to be configured on NSX and in vCenter Server.

- In vSphere:
  - Create a VDS switch.
    - Set MTU to at least 1600
    - Add ESXi hosts to the switch. These hosts are later prepared as NSX transport nodes.
    - Assign uplinks to physical NICs.
- In NSX:
  - When configuring a transport node, map uplinks created in NSX uplink profile with uplinks in VDS.

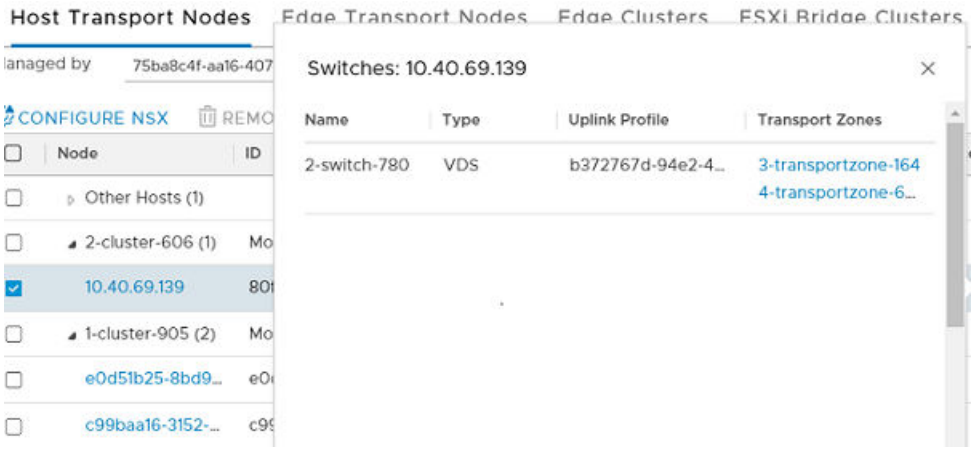
For more details on preparing a host transport node on a VDS switch, see the *NSX Installation Guide*.

The following parameters can only be configured in a VMware vCenter on a VDS backed host switch:

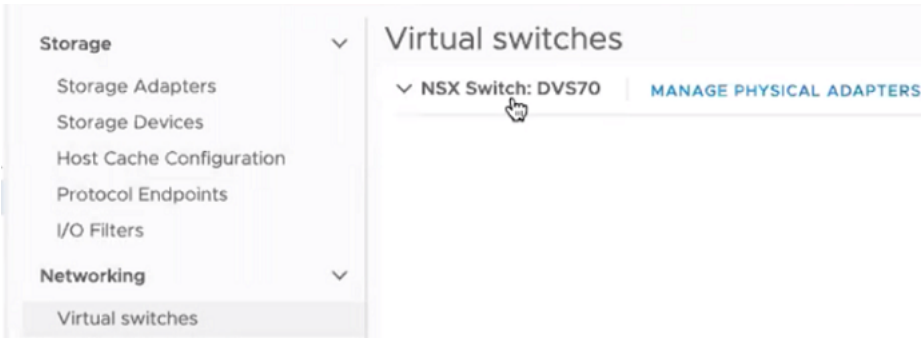
Configuration	VDS	NSX	Description
MTU	In VMware vCenter, set an MTU value on the switch. **Note** A VDS switch must have an MTU of 1600 or higher. In VMware vCenter, select VDS, click **Actions → Settings → Edit Settings**.	Any MTU value set in an NSX uplink profile is overriden.	As a host transport node that is prepared using VDS as the host switch, the MTU value needs to be set on the VDS switch in vCenter Server.
Uplinks/LAGs	In VMware vCenter, configure Uplinks/LAGs on a VDS switch. In VMware vCenter, select VDS, click **Actions → Settings → Edit Settings**.	When a transport node is prepared, the teaming policy on NSX is mapped to uplinks/LAGs configured on a VDS switch.	As a host transport node that is prepared using VDS as the host switch, the uplink or LAG are configured on the VDS switch. During configuration, NSX requires teaming policy be configured for the transport node. This teaming policy is mapped to the uplinks/LAGs configured on the VDS switch.
NIOC	Configure in VMware vCenter. In VMware vCenter, select VDS, click **Actions → Settings → Edit Settings**.	NIOC configuration is not available when a host transport node is prepared using a VDS switch.	As a host transport node that is prepared using VDS as the host switch, the NIOC profile can only be configured in vCenter Server.
Link Layer Discovery Protocol (LLDP)	Configure in VMware vCenter. In VMware vCenter, select VDS, click **Actions → Settings → Edit Settings**.	LLDP configuration is not available when a host transport node is prepared using a VDS switch.	As a host transport node that is prepared using VDS as the host switch, the LLDP profile can only be configured in vCenter Server.
Add or Manage Hosts	Manage in VMware vCenter. In VMware vCenter, go to **Networking → VDS Switch → Add and Manage Host..**	Prepared as transport nodes in NSX.	Before preparing a transport node using a VDS switch, that node must be added to the VDS switch in vCenter Server.

Note NIOC profiles, Link Layer Discovery Protocol (LLDP) profile, and Link Aggregation Group (LAG) for these virtual machines are managed by VDS switches and not by NSX. As a vSphere administrator, configure these parameters from VMware vCenter UI or by calling VDS API commands.

After preparing a host transport node with VDS as a host switch, the host switch type displays VDS as the host switch. It displays the configured uplink profile in NSX and the associated transport zones.



In VMware vCenter, the VDS switch used to prepare NSX hosts is created as an NSX Switch.



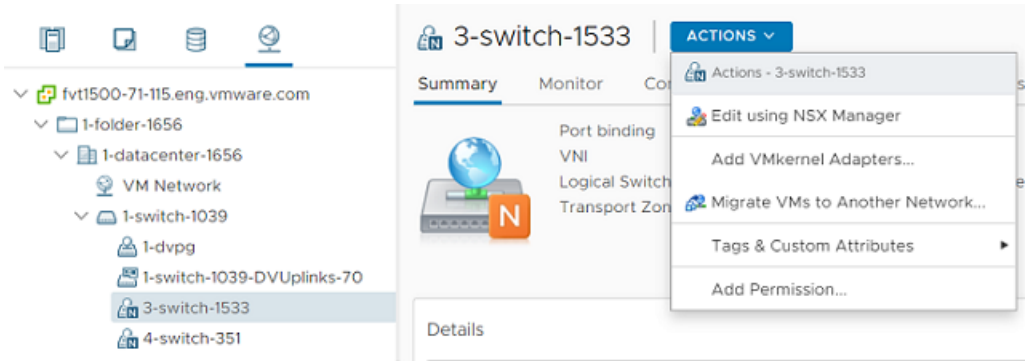
Managing NSX Distributed Virtual Port Groups

A transport node prepared with VDS as a host switch ensures that segments created in NSX is realized as an NSX Distributed Virtual port group on a VDS switch and Segment in NSX .

In earlier versions of NSX, a segment created in NSX are represented as an opaque network in vCenter Server. When running NSX on a VDS switch, a segment is represented as an NSX Distributed Virtual Port Groups.

Any changes to the segments on the NSX network are synchronized in VMware vCenter.

In vCenter Server, an NSX Distributed Virtual Port Group is represented as .



Any NSX segment created in NSX is realized in VMware vCenter as an NSX object. A VMware vCenter displays the following details related to NSX segments:

- NSX Manager
- Virtual network identifier of the segment
- Transport zone
- Attached virtual machines

The port binding for the segment is by default set to **`Ephemeral`**. Switching parameters for the switch that are set in NSX cannot be edited in VMware vCenter and conversely.

Important In a vCenter Server, an NSX Distributed Virtual port group realized does not require a unique name to differentiate it with other port groups on a VDS switch. So, multiple NSX Distributed Virtual port groups can have the same name. Any vSphere automations that use port group names might result in errors.

In vCenter Server, you can perform these actions on an NSX Distributed Virtual Port Group:

- Add VMkernel Adapters.
- Migrate VMs to Another Network.

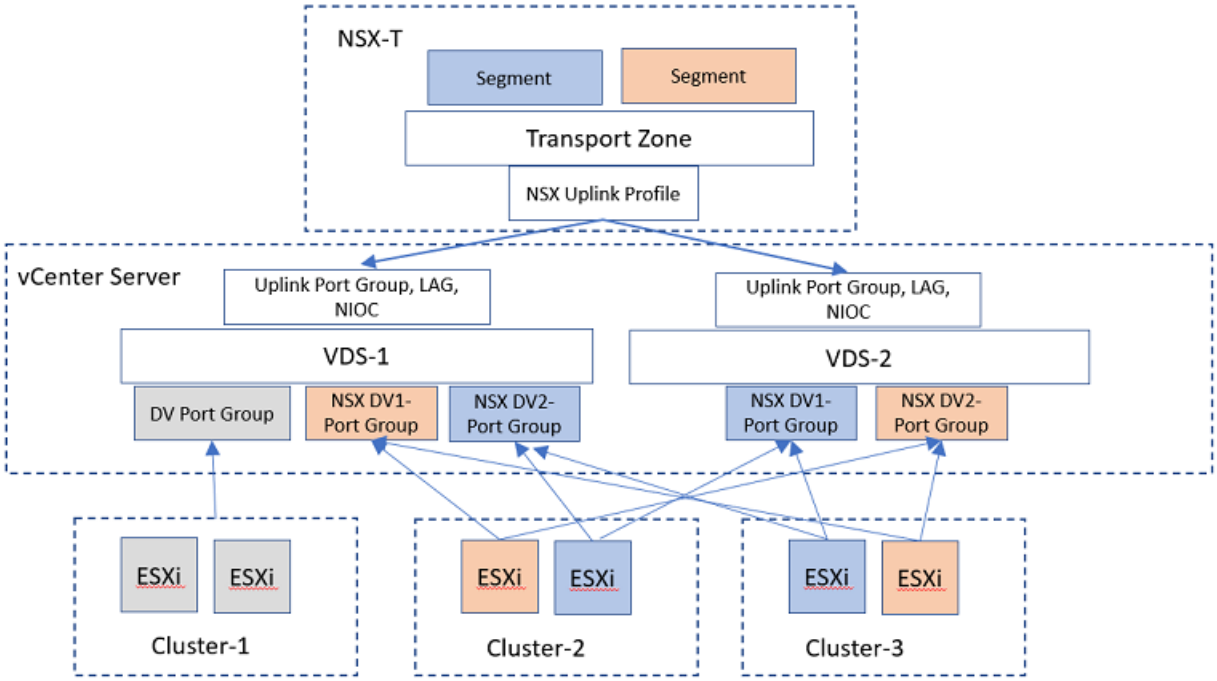
However, NSX objects related to an NSX Distributed Virtual port group can only be edited in NSX Manager. You can edit these segment properties:

- Replication Mode for the segment
- VLAN trunk ID used by the segment
- Switching Profiles (for example, Port Mirroring)
- Ports created on the segment

For details on configuring a vSphere Distributed Virtual port group, refer to the *vSphere Networking Guide*.

NSX Cluster Prepared with VDS

An example of an NSX cluster prepared using VDS as the host switch.



In the sample topology diagram, two VDS switches are configured to manage NSX traffic and vSphere traffic.

VDS-1 and VDS-2 are configured to manage networking for ESXi hosts from Cluster-1, Cluster-2, and Cluster-3. Cluster-1 is prepared to run only vSphere traffic, whereas, Cluster-2 and Cluster-3 are prepared as host transport nodes with these VDS switches.

In vCenter Server, uplink port groups on VDS switches are assigned physical NICs. In the topology, uplinks on `VDS-1` and `VDS-2` are assigned to physical NICs. Depending on the hardware configuration of the ESXi host, you might want to plan out how many physical NICs to be assigned to the switch. In addition to assiging uplinks to the VDS switch, MTU, NIOC, LLDP, LAG profiles are configured on the VDS switches.

After VDS switches are configured in NSX, add an uplink profile.

When preparing a cluster by applying a transport node profile (on a VDS switch), the uplinks from the transport node profile is mapped to VDS uplinks.

After preparing the clusters, ESXi hosts on cluster-2 and Cluster-3 manage NSXtraffic, while cluster-1 manage vSphere traffic.

APIs to Configure vSphere Distributed Switch on NSX

NSX API commands to support vSphere Distributed Switch on NSX.

API Changes for vSphere Distributed Switch

For detailed information related to API calls, see the *NSX API Guide*.

Note Configuration done using API commands is also possible from the VMware vCenter user interface. For more information on creating a NSX transport node using Sphere Distributed Switch as host switch, refer to the *Configure a Managed Host Transport Node* topic in the *NSX Installation Guide*.

API	NSX on vSphere Distributed Switch (VDS)

Create a Transport node for a Discovered node.	`/api/v1/fabric/discovered-nodes/<external-id/discovered-node-id>?action=cr`
	{
	"node_id":
	"d7ef478b-752c-400a-b5f0-207c04567e5d", "host_switch_spec": {
	"host_switches": [
	{
	"host_switch_name": "vds-1",
	"host_switch_id":
	"50 2b 92 54 e0 80 d8 d1-ee ab 8d a6 7b fd f9 4b",
	"host_switch_type": "VDS",
	"host_switch_mode": "STANDARD",
	"host_switch_profile_ids": [
	{
	"key": "UplinkHostSwitchProfile",
	"value":
	"159353ae-c572-4aca-9469-9582480a7467"
	}],
	"pnics": [],
	"uplinks": [
	{
	"vds_uplink_name": "Uplink 2",
	"uplink_name": "nsxuplink1"
	}],
	"is_migrate_pnics": false,
	"ip_assignment_spec": {
	"resource_type": "AssignedByDhcp"
	},
	"cpu_config": [],
	"transport_zone_endpoints": [
	{
	"transport_zone_id":
	"06ba5326-67ac-4f2c-9953-a8c5d326b51e",
	"transport_zone_profile_ids": [
	{
	"resource_type": "BfdHealthMonitoringProfile",
	"profile_id":
	"52035bb3-ab02-4a08-9884-18631312e50a"
	}] }],
	"vmk_install_migration": [],
	"pnics_uninstall_migration": [],
	"vmk_uninstall_migration": [],
	"not_ready": false
	}
	],
	"resource_type": "StandardHostSwitchSpec"
	},
	"transport_zone_endpoints": [],
	"maintenance_mode": "DISABLED",
	"is_overridden": false,
	"resource_type": "TransportNode",
	"display_name": "TestTN",
	}

VM configuration	`vim.vm.device.VirtualEthernetCard.DistributedVirtualPortBackingInfo`

API	NSX on vSphere Distributed Switch (VDS)
VMkernel NIC	`vim.dvs.DistributedVirtualPort`
Physical NIC to Uplink Mapping	API: `vim.host.NetworkSystem:networkSystem.updateNetworkConfig` Property: `vim.host.NetworkConfig.proxySwitch`
MTU	API: `vim.dvs.VmwareDistributedVirtualSwitch.reconfigure` Property: `VmwareDistributedVirtualSwitch.ConfigSpec.maxMtu`
LAG	API: `vim.dvs.VmwareDistributedVirtualSwitch.updateLacpGroupConfig` Property: `vim.dvs.VmwareDistributedVirtualSwitch.LacpGroupSpec`
NIOC	API:`vim.dvs.VmwareDistributedVirtualSwitch.reconfigure` Property: `vim.dvs.VmwareDistributedVirtualSwitch.ConfigSpec.infrastructureT`
LLDP	API: `vim.dvs.VmwareDistributedVirtualSwitch.reconfigure` Property: `vim.dvs.VmwareDistributedVirtualSwitch.ConfigSpec.linkDiscoveryPr`

Feature Support in a vSphere Distributed Switch Enabled to Support NSX

Comparison of features supported by a VDS switch version earlier to 7.0 and VDS version 7.0 or later (NSX enabled).

IPFIX and Port Mirroring

An NSX transport node prepared with a VDS switch supports IPFIX, port mirroring.

See Port Mirroring on a vSphere Distributed Switch.

See IPFIX Monitoring on a vSphere Distributed Switch.

SR-IOV support

SR-IOV is supported on a vSphere Distributed Switch but not on a NSX Virtual Distributed Switch.

Feature	NSX Virtual Distributed Switch	vSphere Distributed Switch
SR-IOV	No	Yes (vSphere 7.0 and later)

Stateless Cluster Host Profile Support

Feature	NSX Virtual Distributed Switch	vSphere Distributed Switch
Host Profile Stateless	Yes	Yes (vSphere 7.0 and later) No (when VMkernel adapters are connected to NSX Port Group on vSphere Distributed Switch.

Distributed Resource Scheduler Support

Source Host	Destination Host	DRS (NIOC Configured)	vSphere
vSphere Distributed Switch-A	vSphere Distributed Switch-B	No	No
vSphere Distributed Switch-A	vSphere Distributed Switch-A	Yes	7.0

vMotion Support

vMotion between source vSphere Distributed Switch and destination vSphere Distributed Switch. Both VDS switches are enabled to support NSX.

Source / VDS	Destination / VDS	Compute vMotion	Storage vMotion
vSphere Distributed Switch-A (VMware vCenter -A)	vSphere Distributed Switch-A (VMware vCenter-A)	Yes	Yes
vSphere Distributed Switch-A (VMware vCenter -A)	vSphere Distributed Switch-B (VMware vCenter -A)	Yes	Yes
vSphere Distributed Switch-A (VMware vCenter -A)	vSphere Distributed Switch-B (VMware vCenter -B)	Yes	Yes
Segment-A (VMware vCenter -A)	Segment-B (VMware vCenter-A)	No	No
Segment-A (VMware vCenter -A)	Segment-B (VMware vCenter -B)	No	No
Transport Zone-A	Transport Zone-B	No	No
NSX-A	NSX-B	No	No

vMotion between vSphere Distributed Switch (NSX enabled) and NSX Virtual Distributed Switch

Source / VDS	Destination / NSX Virtual Distributed Switch	Compute vMotion	Storage vMotion
VMware vCenter-A	VMware vCenter-A	Yes	Yes
VMware vCenter-A	VMware vCenter-B	Yes	Yes
Segment-A (VMware vCenter-A)	Segment-B (VMware vCenter-A)	No	No
Segment-A (VMware vCenter-A)	Segment-B (VMware vCenter-B)	No	No

Source / VDS	Destination / NSX Virtual Distributed Switch	Compute vMotion	Storage vMotion
Transport Zone-A	Transport Zone-B	No	No
NSX-A	NSX-B	No	No

vMotion between vSphere Distributed Switch (NSX enabled) and vSphere Standard Switch or vSphere Distributed Switch

Source / VDS	Destination / NSX Virtual Distributed Switch	Compute vMotion	Storage vMotion
VMware vCenter-A	VMware vCenter-A	Yes	Yes
VMware vCenter-A	VMware vCenter-B	Yes	Yes
Segment-A (VMware vCenter-A)	Segment-B (VMware vCenter-A)	No	No
Segment-A (VMware vCenter-A)	Segment-B (VMware vCenter-B)	No	No
Transport Zone-A	Transport Zone-B	No	No
NSX-A	NSX-B	No	No

Enhanced Networking Stack

Both VDS and NSX Virtual Distributed Switches support all features of the enhanced networking stack.

LACP

- VDS does not support LACP in Active mode.
- NSX Virtual Distributed Switch supports LACP in Active mode.

Scale Supported in vSphere 7.0

Parameter	NSX Virtual Distributed Switch
Logical Switch	■ NSX Distributed Virtual port groups (in VMware vCenter) support 10000 X N, where N is the number of VDS switches in vCenter Server. ■ NSX supports 10000 segments.

Relationship between NSX Distributed Virtual port groups and Hostd memory on the host.

NSX Distributed Virtual Port Groups	Minimum Hostd Memory	Supported VMs
5000	600 MB	191
10000	1000 MB	409
15000	1500 MB	682

License for vSphere Distributed Switch

For earlier versions of NSX, a vSphere Enterprise Plus license is required for the vSphere Distributed Switch 7.0 feature. Starting NSX 3.1.1, the NSX Data Center and NSX Firewall licenses support the use of vSphere Distributed Switch 7.0 for all editions of VMware vCenter and vSphere.

Note With NSX licenses, you get an equivalent number of CPU licenses to use the vSphere Distributed Switch feature. However, the NSX licenses do not provide an upgrade to vSphere Enterprise Plus. The NSX licenses only apply to VDS on the hosts where NSX is deployed.

Procedure

1 Add a VMware vCenter compute manager if you do not have one already registered with NSX. See Add a Compute Manager.

   For a VMware vCenter compute manager that has already been registered, edit the compute manager and provide your credentials for reauthentication.

   Note To use NSX Data Center and NSX Firewall licenses for the vSphere Distributed Switch 7.0 feature, the VMware vCenter user must either be an administrator, or the user must have *Global.Licenses* privileges and be a member of the *LicenseService.Administrators* group.

2 Log in to VMware vCenter and verify the NSX for vSphere solution asset exists. You can use the NSX for vSphere solution asset for NSX deployments.

3 Assign your NSX Data Center or NSX Firewall license to the license asset in VMware vCenter.

   Note The license asset in VMware vCenter is assigned the default NSX for vShield Endpoint license. To use vSphere Distributed Switch, you need any valid NSX license other than the default NSX for vShield Endpoint license.

Results

The vSphere Distributed Switch 7.0 feature is now available and you can attach hosts to a vSphere Distributed Switch.

Enhanced Datapath

Enhanced Datapath is a networking stack mode that provides superior network performance. It is primarily targeted for NFV workloads that require higher performance than regular workloads.

On an ESXi host, configure a vSphere Distributed Switch in NSX in the Enhanced Datapath mode. In the Enhanced Datapath mode, you can configure overlay traffic and VLAN traffic.

Automatically Assign ENS Logical Cores

Automatically assign logical cores to vNICs such that dedicated logical cores manage the incoming traffic to and outgoing traffic from vNICs.

With a switch configured in the enhanced datapath mode, if a single logical core is associated to a vNIC, then that logical core processes bidirectional traffic coming into or going out of a vNIC. When multiple logical cores are configured, the host automatically determines which logical core must process a vNIC's traffic.

Assign logical cores to vNICs based on one of these parameters.

- vNIC-count: Host assumes transmission of incoming or outgoing traffic for a vNIC direction requires same amount of the CPU resource. Each logical core is assigned the same number of vNICs based on the available pool of logical cores. It is the default mode. The vNIC-count mode is reliable, but is not optimal for an asymmetric traffic.
- CPU-usage: Host predicts the CPU usage to transmit incoming or outgoing traffic at each vNIC direction by using internal statistics. Based on the usage of CPU to transmit traffic, host changes the logical core assignments to balance load among logical cores. The CPU usage mode is more optimal than vNIC-count, but unreliable when traffic is not steady.

In CPU usage mode, if the traffic transmitted changes frequently, then the predicted CPU resources required and vNIC assignment might also change frequently. Too frequent assignment changes might cause packet drops.

If the traffic patterns are symmetric among vNICs, the vNIC-count option provides reliable behavior, which is less vulnerable to frequent changes. However, if the traffic patterns are asymmetric, vNIC-count might result in packet drops since it does not distinguish the traffic difference among vNICs.

In vNIC-count mode, it is recommended to configure an appropriate number of logical cores so that each logical core is assigned to the same number of vNICs. If the number of vNICs associated to each logical core is different, CPU assignment is unfair and performance is not deterministic.

When you connect or remove a vNIC or a logical core, the host automatically reflects the changes.

Procedure

- To switch from one mode to another mode, run the following command.

  ```
  set ens lcore-assignment-mode <host-switch-name> <ens-lc-mode>
  ```

  Where, *<ens-lc-mode>* can be set to the mode **`vNIC-count`** or **`cpu-usage`**.

  `vNIC-count` is vNIC/Direction count-based logical core assignment.

  `cpu-usage` is CPU usage-based logical core assignment.

Configure Guest Inter-VLAN Routing

On overlay networks, NSX supports routing of inter-VLAN traffic on an L3 domain. During routing, virtual distributed router (VDR) uses VLAN ID to route packets between VLAN subnets.

Inter-VLAN routing overcomes the limitation of 10 vNICs that can be used per VM. NSX supporting inter-VLAN routing ensures that many VLAN subinterfaces can be created on the vNIC and consumed for different networking services. For example, one vNIC of a VM can be divided into several subinterfaces. Each subinterface belongs to a subnet, which can host a networking service such as SNMP or DHCP. With Inter-VLAN routing, for example, a subinterface on VLAN-10 can reach a subinterface on VLAN-10 or any other VLAN.

Each vNIC on a VM is connected to a switch through the parent logical port, which manages untagged packets.

To create a subinterface, on a switch configured in Enhanced mode, create a child port using the API with an associated VIF using the API call described in the procedure. The subinterface tagged with a VLAN ID is associated to a new logical switch, for example, VLAN10 is attached to logical switch LS-VLAN-10. All subinterfaces of VLAN10 have to be attached to LS-VLAN-10. This 1–1 mapping between the VLAN ID of the subinterface and its associated logical switch is an important prerequisite. For example, adding a child port with VLAN20 to logical switch LS-VLAN-10 mapped to VLAN-10 makes routing of packets between VLANs non-functional. Such configuration errors make the inter-VLAN routing non-functional.

Starting from NSX 3.2.2, logical port proton APIs are replaced with the corresponding segment port policy APIs.

Prerequisites

- Before you associate a VLAN subinterface to a logical switch, ensure that the logical switch does not have any other associations with another VLAN subinterface. If there is a mismatch, inter-VLAN routing on overlay networks might not work.
- Ensure that hosts run ESXi v 6.7 U2 or later versions.

Procedure

1 To create subinterfaces for a vNIC, ensure that the vNIC is updated to a parent port. Make the following REST API call:

```
PATCH https://<nsx-mgr-ip>/policy/api/v1/infra/segments/<Segment to which vNIC is
connected>/ports/<Seg-Port-vNIC>
{
   "attachment": {
     "id": "<Attachment UUID of the vNIC>",
     "type": "PARENT"
  },
   "admin_state": "UP",
   "resource_type": "SegmentPort",
   "display_name": "parentport"
}
}
```

If the logical switch does not have a corresponding segment, you can make the following REST API calls (logical port proton API):

```
PUT https://<nsx-mgr-ip>/api/v1/logical-ports/<Logical-Port UUID-of-the-vNIC>
{
  "resource_type" : "LogicalPort",
  "display_name" : "parentport",
  "attachment" : {
    "attachment_type" : "VIF",
    "context" : {
      "resource_type" : "VifAttachmentContext",
      "vif_type": "PARENT"
    },
    "id" : "<Attachment UUID of the vNIC>"
  },
  "admin_state" : "UP",
  "logical_switch_id" : "UUID of Logical Switch to which the vNIC is connected",
  "_revision" : 0
}
```

2 To create child ports for a parent vNIC port on the N-VDS that is associated to the subinterfaces on a VM, make the following API call:

Note Before making the API call, verify that a segment exists to connect child ports with the subinterfaces on the VM.

```
PUT https://<nsx-mgr-ip>/policy/api/v1/infra/segments/<Segment to which child port is
connected>/ports/<Child-port>
{
   "attachment": {
     "id": "<Attachment UUID of the CHILD port>",
     "type": "CHILD",
     "context_id": "<Attachment UUID of the PARENT port from Step 1>",
     "traffic_tag": <VLAN ID>,
     "app_id": "<ID of the attachment>", ==> display id(can be any string). Must be unique.
  },
   "address_bindings": [
    {
       "ip_address": "<IP address to the corresponding VLAN>",
       "mac_address": "<vNIC MAC Address>",
       "vlan_id": <VLAN ID>
    }
  ],
   "admin_state": "UP",
   "resource_type": "SegmentPort",
   "display_name": "<Name of the Child PORT>"
}
```

If the logical switch does not have a corresponding segment, you can make the following REST API calls (logical port proton API):

```
POST https://<nsx-mgr-ip>/api/v1/logical-ports/
{
```

```
  "resource_type" : "LogicalPort",
  "display_name" : "<Name of the Child PORT>",
  "attachment" : {
    "attachment_type" : "VIF",
    "context" : {
      "resource_type" : "VifAttachmentContext",
      "parent_vif_id" : "<UUID of the PARENT port from Step 1>",
      "traffic_tag" : <VLAN ID>,
      "app_id" : "<ID of the attachment>",  ==> display id(can give any string). Must be
unique.
      "vif_type" : "CHILD"
    },
   "id" : "<ID of the CHILD port>"
  },

  "logical_switch_id" : "<UUID of the Logical switch(not the PARENT PORT's logical switch)
to which Child port would be connected to>",
  "address_bindings" : [ { "mac_address" : "<vNIC MAC address>", "ip_address" : "<IP
address to the corresponding VLAN>", "vlan" : <VLAN ID> } ],
  "admin_state" : "UP"
}
```

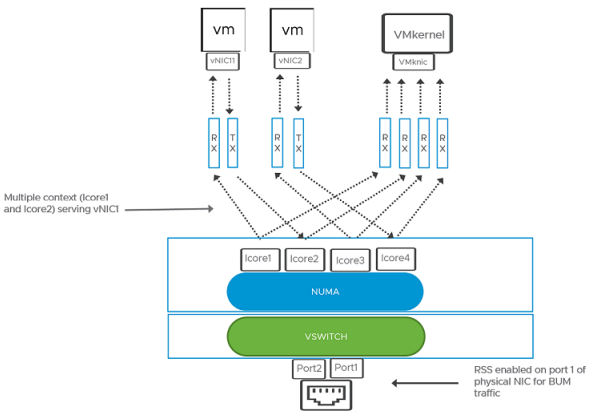
Results

NSX creates subinterfaces on VMs.

Receive Side Scaling

Receive Side Scaling allows multiple cores on the receive side for processing incoming traffic.

Without RSS, receiving ESXi hosts only use one physical queue and hence one lcore for packet processing. When the receive side data increases it creates a bottleneck at the single core. The overall throughput performance might decrease. With RSS enabled on the NIC, you can configure multiple hardware queues to process requests from VMs. Before you use a NIC card to leverage the RSS functionality, use the VMware Compatibility Guide for I/O to confirm whether the NIC card driver supports RSS. Most of the NIC cards support at least 4 queues. So, RSS might provide 4x throughput performance improvements.

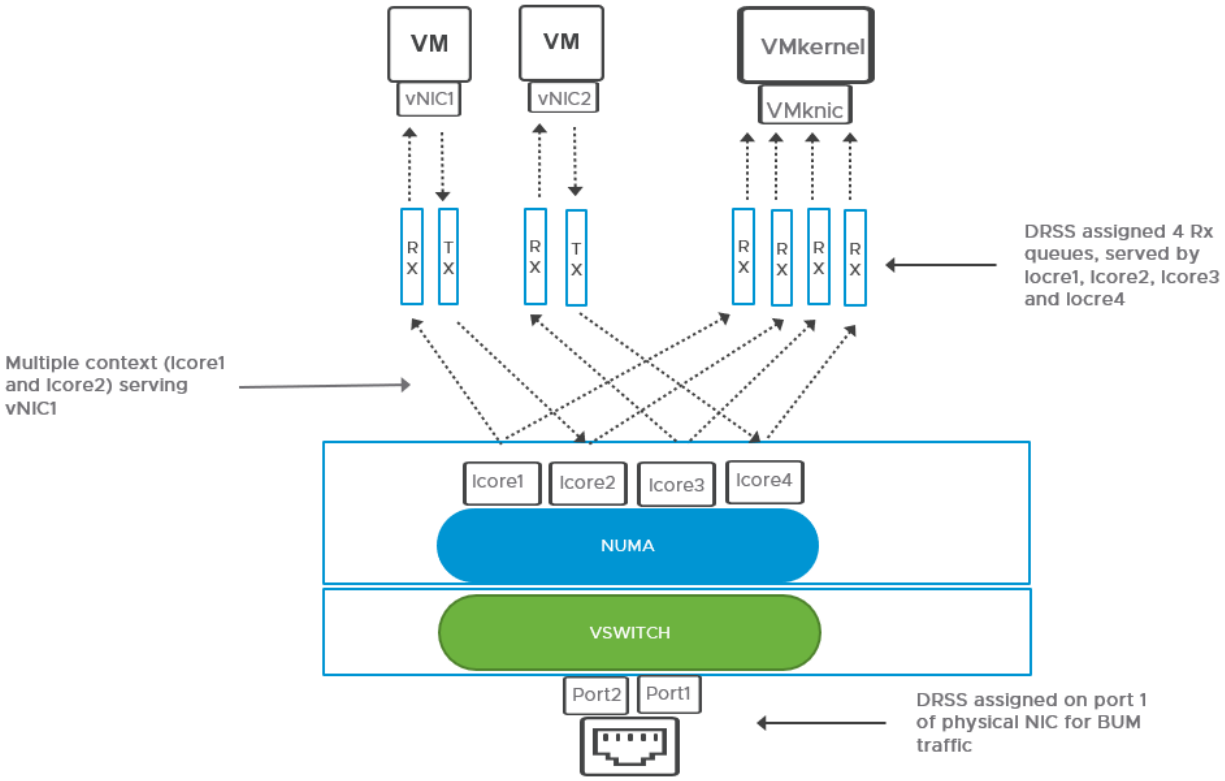


You can choose to configure RSS in these modes:

- enable Default Queue Receive Side Scaling (DRSS)
- enable RSS engine dedicated to a single vNIC queue
- enable RSS engine shared by multiple vNIC queues

Configure Multiple Context on Host Switch

Provide multiple cores to vNICs by configuring the Multiple Context functionality on a host switch running in Enhanced Datapath mode. It helps improve packet performance.



On a host switch configured to run in the **Enhanced Datapath** mode, you can configure Multiple Context functionality for vNIC traffic. Multiple Context means that multiple logical cores can serve Tx (transmit) and Rx (receive) queues, in contrast to the single context, where only one logical core serves both the Tx queue and Rx queue. A Tx and Rx queue pair represents a vNIC queue.

As an admin, you can assign Multiple Context to vNIC queues based on the network traffic load. As traffic load increases for a vNIC queue, a single context or logical core for a specific vNIC queue can prove to be insufficient to load balance traffic. Assigning Multiple Context to that vNIC allocates more vCPU resources to load balance traffic.

As you design for an optimized network and increased throughput, consider these points:

- The number of logical cores assigned depends on the capacity of the host.

- The number of Default Queue RSS (DRSS) configurable on a host depends on the maximum number of physical CPUs available on the host.
- Logical cores can be shared across DRSS and Multiple Context queues.
- Both DRSS and Multiple Context can function independently. However, configuring them together provides additional performance benefits to physical hardware queues (DRSS) and vNIC queues. See Configure Default Queue Receive Side Scaling for more details on configuring DRSS.

Prerequisites

- To configure the Multiple Context functionality for a vNIC, create multiple logical cores on the host.
- Ensure that the host switch is configured in **ENS Interrupt** mode or **Enhanced Datapath** mode. The Multiple Context functionality is not available in the **Standard** mode.

Procedure

1. To verify that host switch is configured to run in **Enhanced Datapath** mode:
   a. Navigate the UI based on the NSX version and select a host:
      - **System > Fabric > Hosts** and select the **Cluster** tab.
      - **System > Fabric > Hosts** and select the **Other Nodes** tab.
   b. Select the transport node.
   c. Select the **Overview** tab and verify `Enhanced Datapath Capable` parameter is set to `Yes`.
2. To configure Multiple Context functionality for vNIC traffic managed through Enhanced Datapath mode, edit configuration options of VMs and set the following parameter value. See the latest *vSphere Virtual Machine Administration guide* for details on how to edit VM configuration options.

   ```
   ethernetX.ctxPerDev = "3"
   ```

   Where, the value `3` indicates that Multiple Context functionality is applied per vNIC queue.

   Other supported values for contexts are:

   - `ethernetX.ctxPerDev =1` indicates that Multiple Context functionality is applied per VM.
   - `ethernetX.ctxPerDev =2` indicates that Multiple Context functionality is applied per vNIC.

Results

Enhanced Datapath improves packet throughput by using the Multiple Context functionality set for vNIC queues.

Configure Default Queue Receive Side Scaling

Improve packet throughput by enabling Default Queue Receive Side Scaling (DRSS) on the NIC card.

After you enable the Default Queue Receive Side Scaling (DRSS) configuration on a NIC port, Enhanced Network Stack (ENS) manages the receive-side data arriving at physical NIC cards. A single port on the physical NIC card makes multiple hardware queues available to receive-side data. Each queue is assigned a local logical core from the non-uniform memory access (NUMA) node. When inbound packets - multicast, unknown, or broadcast - arrive at a physical NIC port, they are distributed across several hardware queues, depending on the availability of logical cores. DRSS reduces bottlenecks processed by a single queue. DRSS is intended to serve broadcast, unknown, or multicast (BUM) traffic.

For example, on a physical NIC card that has two ports, you can configure one port to make multiple hardware queues available to efficiently manage receive-side (Rx) traffic. It can be done by passing DRSS=4,0 value in the ESXi system parameters command. This parameter enables the first physical NIC port for DRSS.

Note If multiple context is not enabled, then configuring vNICs for multiple context does not work.

Prerequisites

- Ensure the NIC card supports Default Queue Receive Side Scaling.

Procedure

1. Install i40en ENS driver NIC driver.
2. If the NIC has two ports, enable RSS on the first port of the physical NIC, by running the command.

   ```
   esxcli system module parameters set -m -i40en_ens -p DRSS=4,0
   ```

   Where, DRSS is enabled for 4 Rx queues on the first port and it is not enabled for Tx queues.

   The number of DRSS queues assigned depends on the number of physical CPUs available on the host.

   Note Depending on the version of the NIC card, by default, the DRSS might be enabled or disabled.

3. If NIC teaming is in use, then configuration of both NIC ports must be the same.

   ```
   esxcli system module parameters set -m -i40en_ens -p DRSS=4,4
   ```

4. Unload Load the NIC driver for module parameters to take effect.
5. Load the NIC driver.

What to do next

Configure multiple context so that ENS module can improve packet throughput of vNIC queues.

Configure NetQ Receive Side Scaling

Enable NetQ Receive Side Scaling to enable vNIC requests to be offloaded to a physical NIC. It improves packet performance of the receive-side data.

When a physical NIC card sends packets to a host, the Enhanced Network Stack (ENS), which runs as the host switch is configured in **Enhanced Datapath** mode, on that host distributes data across different logical cores on NUMA nodes. There are a couple of ways to configure RSS engines.

As a network admin wanting to improve the throughput packet performance of receive-side data, you might want to consider one of these ways to configure RSS to leverage the benefits.

These two modes are:

- RSS engine is dedicated to a single vNIC queue: A dedicated RSS engine completely offloads any request coming from a vNIC to the physical NIC. In this mode, a single RSS engine is dedicated to a single vNIC queue. It improves throughput performance as pNIC manages the recieve side data and shares it among the available hardware queues to serve the request. The vNIC queues are co-located on the same logical core or fastpath as pnic queues.
- RSS engine is shared by multiple vNIC queues: In this mode, multiple hardware queues are made available to vNIC queues. However, the vNIC handling flows might not be aligned with the physical hardware queue that will process data. It means, there is no guarantee that vNIC and physical NICs will be aligned.

Note If Default Queue Receive Side Scaling (DRSS) is enabled on the NIC card, deactivate it.

Prerequisites

- Hosts must be running ESXi version 7 update 3 or later.
- Ensure NIC card supports RSS functionality.
- EDP NETQ RSS is supported from NSX 4.0 and ESXi version 8.0 onwards. Supported inbox drivers are Intel40en (async driver) and Mellanox nmlx. Refer to the driver documentation to confirm whether it has ENS compatible RSS implementation.

Procedure

1. To enable NetQ/RSS, `esxcli system module parameters set -m -i40en_ens -p DRSS=0,0 RSS=1,0`.

   Where, DRSS=0,0 indicates DRSS is deactivated on both NIC ports.

   RSS=1,0 indicates NetQ RSS is enabled on one of the NIC ports.
2. To unload driver, run `vmkload_mod -u i40en_ens`.
3. To reload driver for the RSS setting to take effect, run `vmkload_mod i40en_ens`.

4 Stop the device manager to trigger PCI fastconnect so that it can scan devices and associate the driver with a NIC.

   Run `kill -HUP 'ps | grep mgr | awk '{print $1}'`.

5 To configure multiple RSS engines to be available to serve RSS requests from vNICs, configure these parameters in the `.vmx` file of VM.

   `ethernet.pnicfeatures = '4'`, which indicates RSS feature is requested by vNICs.

   `ethernet.ctxPerDev = '3'`, which indicates that multiple contexts (multiple logical cores) are enabled to process each vNIC. The VMs connected to the vSphere switch are configured for multiple queues. It means multiple logical cores of a NUMA node can process the Tx and Rx traffic coming from vNICs.

   When multiple vNICs request RSS offloading, the Enhanced Network Stack (ENS) does not offload their RSS requests to the pnic, but the shared RSS engine processes their requests. For shared RSS, multiple RSS queues are available but co-location of a vNIC queue or a pNIC queue is not guranteed.

6 To configure a dedicated RSS engine to process requests from a vNIC, configure these parameters in the `.vmx` file of the VM.

   `ethernet.rssoffload=True`,

   With the preceding configuration enabled, RSS requests from a vNIC is offloaded to the physical NIC. Only one vNIC can offload its requests to an RSS engine. In this mode, vNIC queues are aligned to the pNIC queues.

7 Verify that packet flow is distributed on the hardware queues provided by the RSS engine.

   Run the following commands.

   ```
   vsish
   ```

   ```
   get /net/pNics/vmnicX/stats
   ```

   Sample output:

```
rxq0: pkts=0 bytes=0 toFill=2047 toProc=0 noBuf=0 csumErr=0
rxq1: pkts=0 bytes=0 toFill=2047 toProc=0 noBuf=0 csumErr=0
rxq2: pkts=0 bytes=0 toFill=2047 toProc=0 noBuf=0 csumErr=0
rxq3: pkts=0 bytes=0 toFill=2047 toProc=0 noBuf=0 csumErr=0
rxq4: pkts=0 bytes=0 toFill=2047 toProc=0 noBuf=0 csumErr=0
rxq5: pkts=0 bytes=0 toFill=2047 toProc=0 noBuf=0 csumErr=0
rxq6: pkts=0 bytes=0 toFill=2047 toProc=0 noBuf=0 csumErr=0
rxq7: pkts=0 bytes=0 toFill=2047 toProc=0 noBuf=0 csumErr=0
txq0: pkts=0 bytes=0 toFill=0 toProc=0 dropped=0
txq1: pkts=0 bytes=0 toFill=0 toProc=0 dropped=0
txq2: pkts=0 bytes=0 toFill=0 toProc=0 dropped=0
txq3: pkts=0 bytes=0 toFill=0 toProc=0 dropped=0
```

```
txq4: pkts=0 bytes=0 toFill=0 toProc=0 dropped=0
txq5: pkts=0 bytes=0 toFill=0 toProc=0 dropped=0
txq6: pkts=0 bytes=0 toFill=0 toProc=0 dropped=0
txq7: pkts=0 bytes=0 toFill=0 toProc=0 dropped=0
```

Configure VTEP HA

To avoid connectivity loss of workloads (VMs or Containers) associated to a virtual tunnel endpoint (VTEP), configure VTEP HA host switch profile and apply it to cluster transport nodes or individual transport nodes.

As overlay networks spanning across L2 or L3 networks use encapsulation protocols, packets egressing or ingressing from workloads (VMs or Containers) are encapsulated within the outer packet of a Virtual Tunnel End Point (VTEP). If a VTEP goes into a faulty state, the VMs associated with it are not reachable. There is a one-to-one mapping between VMs to VTEPs.

A VTEP is faulty when one of the following points are true:

- All BFD sessions of a VTEP are down due to issues in the underlay fabric.
- VTEP does not have an IP address.

A healthy VTEP has the following characteristics:

- It has an IP address.
- At least one BFD session involving the VTEP is Up.

When NSX discovers a VTEP is down, vNICs of VMs are migrated from the faulty VTEP to a healthy VTEP on the same ESXi host. If a ESXi has two VTEPs and both are faulty, then VM vNICs cannot be migrated.

Each VM can have more than one vNIC. Each vNIC is associated to a single VTEP. The packets from the vNIC are encapsulated before being sent out on the uplink. There can be multiple vNICs mapped to a single VTEP. The mapping is decided based on uplink teaming policies.

Note VTEP HA only supports IPv4 VTEPs.

Prerequisites

At least two VTEPs are configured on a single ESXi host. Ensure the uplink profile you use to create a transport node profile (TNP) or individual transport node configuration uses at least two active uplinks (VTEPs).

Procedure

1 Create a VTEP HA host switch profile.

```
PUT https://<nsx-policy-manager>/policy/api/v1/infra/host-switch-profiles/
vtephaprofile1
```

```
{
"enabled": "true",
```

```
"failover_timeout":"5",
"auto_recovery" : "true",
"auto_recovery_initial_wait" : "300",
"auto_recovery_max_backoff" : "86400",
"resource_type": "PolicyVtepHAHostSwitchProfile",
"display_name": "VtepHAProfile1"
}
```

where,

`enabled`: Indicates VTEP HA is enabled or not. The "true" value indicates VTEP HA is enabled and "false" indicates it is disabled.

`failover_timeout`: Is the VTEP High Availability failover timeout (in seconds). This property controls after how much time should VMs vNIC be moved to an alternate healthy VTEP once fault is detected at a vmknic. Default timeout is 5 seconds. Maximum value supported is up to 60 seconds.

`auto_recovery`: Specifies status of autonomous recovery option for the VTEP High Availability feature. When `auto_recovery` is set to true, faulted vmknic is checked internally after every cyclic sequence. If the faulty vmknic recovers, VM is moved back to its initial VTEP mapping, provided the VMs vNIC was mapped to that VTEP. By default, `auto_recovery` is set to true.

During auto recovery VMs face outage if VTEP is still in faulty state. The outage time is calculated using the formula: `failover_timeout + bfd timeout + 1`. The default `bfd timeout` is `4` seconds.

`auto_recovery_initial_wait`: This property controls after how much time the autonomous recovery should start. The minimum initial wait period is 300 seconds, while the maximum wait period is 3600 seconds. If initial wait period is 300 seconds, then auto recovery starts at 360 seconds.

`auto_recovery_max_backoff`: This property controls the maximum autorecovery time the system supports. The default value is 86400 seconds. The autorecovery cycle time after reaching the maximum wait period is set to the `auto_recovery_max_backoff` to avoid aggressive autorecovery. Autorecovery happens only after reachig the value set for `auto_recovery_max_backoff`. The maximum value supported for this parameter is 86400.

2 To attach a VTEP HA profile to all ESXi hosts, create a TNP with the VTEP HA profile key and value properties added to the Host Switch Profiles section or prepare individual transport nodes with the VTEP HA profile key and value properties.

`PUT https://<nsx-policy-manager>/policy/api/v1/infra/host-transport-node-profiles/<tnp-id>`

In the `<tnp-id>` parameter, enter the transport node profile ID.

```
{
    "host_switch_spec": {
        "host_switches": [
            {
```

```
                "host_switch_name": "vtepha_vds",
                "host_switch_id": "50 22 ee c4 f6 40 79 8b-0e a4 2b da 6a 4c 36 b3",
                "host_switch_type": "VDS",
                "host_switch_mode": "ENS_INTERRUPT",
                "host_switch_profile_ids": [
                    {
                        "key": "UplinkHostSwitchProfile",
                        "value": "/infra/host-switch-profiles/b32e6ce6-f1ba-4e31-
a2c4-33d550505cdd"
                    },
                    {
                        "key": "VtepHAHostSwitchProfile",
                        "value": "/infra/host-switch-profiles/vtephaprofile1"
                    }
                ],
                "uplinks": [
                    {
                        "vds_uplink_name": "Uplink 1",
                        "uplink_name": "uplink-1"
                    },
                    {
                        "vds_uplink_name": "Uplink 2",
                        "uplink_name": "uplink-2"
                    }
                ],
                "is_migrate_pnics": false,
                "ip_assignment_spec": {
                    "ip_pool_id": "/infra/ip-pools/v4",
                    "resource_type": "StaticIpPoolSpec"
                },
                "cpu_config": [],
                "transport_zone_endpoints": [
                    {
                        "transport_zone_id": "/infra/sites/default/enforcement-points/
default/transport-zones/de47a6b9-fa4c-4bf3-bd75-385859895949",
                        "transport_zone_profile_ids": []
                    }
                ],
                "not_ready": false
            }
        ],
        "resource_type": "StandardHostSwitchSpec"
    },
    "ignore_overridden_hosts": false,
    "resource_type": "PolicyHostTransportNodeProfile"
 }
```

Note If you do not configure TNP with VTEP HA, NSX creates a default VTEP HA profile. This default VTEP HA profile is configured with default values. VTEP HA is disabled.

3 To know the VTEP HA profile ID that you created in Step 1, run the following API:

```
get https://<nsx-manager>/policy/api/v1/infra/host-switch-profiles/
vtephaprofile1
```

In the GET response, you can get the value of VTEP HA profile from the `id` field.

4 To verify whether VTEP HA host switch profile is applied to transport nodes, run the following API.

`GET https://<nsx-manager>/policy/api/v1/infra/host-transport-nodes-profiles/<tnp_id>`

```
 "host_switch_profile_ids": [
     {
        "key": "VtepHAHostSwitchProfile",
        "value": "/infra/host-switch-profiles/<vtephaprofile1>"
      }
    ],
```

`<vtephaprofile1>` is the ID that is returned when you created VTEP HA profile in Step 1.

5 You can also trigger a manual recovery of an faulty VTEP and migrate associated VMs to a healthy VTEP without waiting for auto recovery to begin.

`POST https://<nsx-mgr>/policy/api/v1/infra/sites/<site-id>/enforcement-points/<enforcementpoint-id>/host-transport-nodes/<host-transport-node-id>/vteps/actions`

```
{
    "type": "TransportNodeVTEPRecoveryRequest",
    "device_name": "vmk10"
}
```

where vmk10 is the VTEP name.

6 If vmknics of VTEPs are faulty, run the alarm API to view details related to the error or view the error on NSX Manager Dashboard.

NSX might take at least 60 seconds to raise an alarm if a vmknic is faulty. Similarly, when a VTEP comes back up again after being in faulty state, NSX can take up to 60 seconds to reflect the new state.

`GET https://<nsx-manager>/api/v1/alarms?feature_name=tep_health`

7 After fixing either IP or BFD-related errors, verify the `state` parameter to know the vmknics or VTEP state.

`GET https://<nsx-manager>/api/v1/transport-nodes/<node-id>/network/interfaces?source=realtime`

```
{
        "interfaceId": "vmk10",
        "linkStatus": "UP",
        "adminStatus": "UP",
        "mtu": "1600",
        "interfaceAlias": [{
          "broadcastAddress": "133.117.22.255",
          "ipAddress": {
```

```
            "ipv4": 2239043120
          },
          "ipConfiguration": "STATIC",
          "netmask": "255.255.255.0",
          "macAddress": "00:50:56:66:67:a6"
        }],
        "state": "NORMAL"
      },
```

Results

If BFD is down, NSX determines whether faulty VTEP is back UP using auto recovery or manual recovery mechanism.

If IP is unavailable initially but is reassigned to a faulty VTEP, NSX uses an in-built mechanism to determine when VTEP is back UP. It does not rely on manual or auto recovery mechanism.

If the faulty VTEP does not come back UP:

- During auto recovery or manual recovery process, VM vNIC might face a temporary network outage.
- Network outage might be because during BFD check the VM vnic is still mapped to the faulty VTEP.
- Network outage time is calculated using the formula: `failover_timeout + bfd timeout + 1`, which is 10 seconds. The default `bfd timeout` is `4` seconds.

If faulty VTEP still does not come back UP, VM vNIC is remapped to a healthy VTEP that is available on the host. If no VTEPs are healthy, then VM vNIC is mapped to its original VTEP.

VTEP Maintenance Mode

You can put a VTEP in maintenance mode provided there is an alternate healthy VTEP on the host.

Prerequisites

- Ensure at least one healthy VTEP is present on the host.

Procedure

1 To put a VTEP in Administrative Down mode, call the following API:

`POST https://<nsx-mgr>/policy/api/v1/infra/sites/<site-id>/enforcement-points/<enforcementpoint-id>/host-transport-nodes/<host-transport-node-id>/vteps/actions`

```
{
     "type": "TransportNodeVtepAdminStateMgmtRequest",
     "device_name": "vmk10"
     "admin_state_up": <state>
}
```

where,

`state` is **`false`** for configuring VTEP in Administrative Down mode.

`state` is **`true`** for configuring VTEP in non-Administrative Down mode, which is to bring VTEP back up from Administrative Down mode to normal state.

This operation does not persist after restarting the transport node.

2 Move the VM vNIC of the VTEP that is Administrative Down mode to another healthy VTEP.

Results

VM vNIC from VTEP in Administrative Down mode is moved to another healthy VTEP.

Virtual Private Network (VPN)

7

NSX supports IPSec Virtual Private Network (IPSec VPN) and Layer 2 VPN (L2 VPN) on an NSX Edge node. IPSec VPN offers site-to-site connectivity between an NSX Edge node and remote sites. With L2 VPN, you can extend your data center by enabling virtual machines to keep their network connectivity across geographical boundaries while using the same IP address.

Note IPSec VPN and L2 VPN are not supported in the NSX limited export release.

You must have a working NSX Edge node, with at least one configured Tier-0 or Tier-1 gateway, before you can configure a VPN service. For more information, see "NSX Edge Installation" in the *NSX-T Data Center Installation Guide*.

Beginning with NSX 2.4, you can also configure new VPN services using the NSX Manager user interface. In earlier releases of NSX, you can only configure VPN services using REST API calls.

Important When using NSX 2.4 or later to configure VPN services, you must use new objects, such as Tier-0 gateways, that were created using the NSX Manager UI or Policy APIs that are included with NSX 2.4 or later release. To use existing Tier-0 or Tier-1 logical routers that were configured before the NSX 2.4 release, you must continue to use API calls to configure a VPN service.

System-default configuration profiles with predefined values and settings are made available for your use during a VPN service configuration. You can also define new profiles with different settings and select them during the VPN service configuration.

The Intel QuickAssist Technology (QAT) feature on a bare metal server is supported for IPSec VPN bulk cryptography. Support for this feature began with NSX 3.0. For more information on support of the QAT feature on bare metal servers, see the *NSX Installation Guide*.

This chapter includes the following topics:

- Understanding IPSec VPN
- Understanding Layer 2 VPN
- Adding VPN Services
- Adding IPSec VPN Sessions
- Adding L2 VPN Sessions
- Add Local Endpoints

- Adding Profiles
- Add an Autonomous Edge as an L2 VPN Client
- Check the Realized State of an IPSec VPN Session
- Understanding TCP MSS Clamping
- Troubleshooting VPN Problems

Understanding IPSec VPN

Internet Protocol Security (IPSec) VPN secures traffic flowing between two networks connected over a public network through IPSec gateways called endpoints. NSX Edge only supports a tunnel mode that uses IP tunneling with Encapsulating Security Payload (ESP). ESP operates directly on top of IP, using IP protocol number 50.

IPSec VPN uses the IKE protocol to negotiate security parameters. The default UDP port is set to 500. If NAT is detected in the gateway, the port is set to UDP 4500.

NSX Edge supports a policy-based or a route-based IPSec VPN.

Beginning with NSX 2.5, IPSec VPN services are supported on both Tier-0 and Tier-1 gateways. See Add a Tier-0 Gateway or Add a Tier-1 Gateway for more information. The Tier-0 or Tier-1 gateway must be in `Active-Standby` high-availability mode when used for an IPSec VPN service. You can use segments that are connected to either Tier-0 or Tier-1 gateways when configuring an IPSec VPN service.

An IPsec VPN service in NSX uses the gateway-level failover functionality to support a high-availability service at the VPN service level. Tunnels are re-established on failover and VPN configuration data is synchronized. Before NSX 3.0 release, the IPSec VPN state is not synchronized as tunnels are being re-established. Beginning with NSX 3.0 release, the IPSec VPN state is synchronized to the standby NSX Edge node when the current active NSX Edge node fails and the original standby NSX Edge node becomes the new active NSX Edge node without renegotiating the tunnels. This feature is supported for both policy-based and route-based IPSec VPN services.

Pre-shared key mode authentication and IP unicast traffic are supported between the NSX Edge node and remote VPN sites. In addition, certificate authentication is supported beginning with NSX 2.4. Only certificate types signed by one of the following signature hash algorithms are supported.

- SHA256withRSA
- SHA384withRSA
- SHA512withRSA

Using Policy-Based IPSec VPN

Policy-based IPSec VPN requires a VPN policy to be applied to packets to determine which traffic is to be protected by IPSec before being passed through the VPN tunnel.

This type of VPN is considered static because when a local network topology and configuration change, the VPN policy settings must also be updated to accommodate the changes.

When using a policy-based IPSec VPN with NSX, you use IPSec tunnels to connect one or more local subnets behind the NSX Edge node with the peer subnets on the remote VPN site.

You can deploy an NSX Edge node behind a NAT device. In this deployment, the NAT device translates the VPN address of an NSX Edge node to a publicly accessible address facing the Internet. Remote VPN sites use this public address to access the NSX Edge node.

You can place remote VPN sites behind a NAT device as well. You must provide the remote VPN site's public IP address and its ID (either FQDN or IP address) to set up the IPSec tunnel. On both ends, static one-to-one NAT is required for the VPN address.

Note DNAT is not supported on tier-0 or tier-1 gateways where policy-based IPSec VPN are configured.

IPSec VPN can provide a secure communications tunnel between an on-premises network and a network in your cloud software-defined data center (SDDC). For policy-based IPSec VPN, the local and peer networks provided in the session must be configured symmetrically at both endpoints. For example, if the cloud-SDDC has the *local* networks configured as **`X, Y, Z`** subnets and the *peer* network is **`A`**, then the on-premises VPN configuration must have **`A`** as the *local* network and **`X, Y, Z`** as the *peer* network. This case is true even when **`A`** is set to **`ANY (0.0.0.0/0)`**. For example, if the cloud-SDDC policy-based VPN session has the *local* network configured as **`10.1.1.0/24`** and the *peer* network as **`0.0.0.0/0`**, at the on-premises VPN endpoint, the VPN configuration must have **`0.0.0.0/0`** as the local network and **`10.1.1.0/24`** as the *peer* network. If misconfigured, the IPSec VPN tunnel negotiation might fail.

The size of the NSX Edge node determines the maximum number of supported tunnels, as shown in the following table.

Table 7-1. Number of IPSec Tunnels Supported

Edge Node Size	# of IPSec Tunnels Per VPN Session (Policy-Based)	# of Sessions Per VPN Service	# of IPSec Tunnels Per VPN Service (16 tunnels per session)
Small	N/A (POC/Lab Only)	N/A (POC/Lab Only)	N/A (POC/Lab Only)
Medium	128	128	2048
Large	128 (soft limit)	256	4096
Bare Metal	128 (soft limit)	512	6000

Restriction The inherent architecture of policy-based IPSec VPN restricts you from setting up a VPN tunnel redundancy.

For information about configuring a policy-based IPSec VPN, see Add an IPSec VPN Service.

Using Route-Based IPSec VPN

Route-based IPSec VPN provides tunneling on traffic based on the static routes or routes learned dynamically over a special interface called virtual tunnel interface (VTI) using, for example, BGP as the protocol. IPSec secures all the traffic flowing through the VTI.

Note

- OSPF dynamic routing is not supported for routing through IPSec VPN tunnels.
- Dynamic routing for VTI is not supported on VPN that is based on Tier-1 gateways.
- Load balancer over IPSec VPN is not supported for route-based VPN terminated on Tier-1 gateways.

Route-based IPSec VPN is similar to Generic Routing Encapsulation (GRE) over IPSec, with the exception that no additional encapsulation is added to the packet before applying IPSec processing.

In this VPN tunneling approach, VTIs are created on the NSX Edge node. Each VTI is associated with an IPSec tunnel. The encrypted traffic is routed from one site to another site through the VTI interfaces. IPSec processing happens only at the VTI.

VPN Tunnel Redundancy

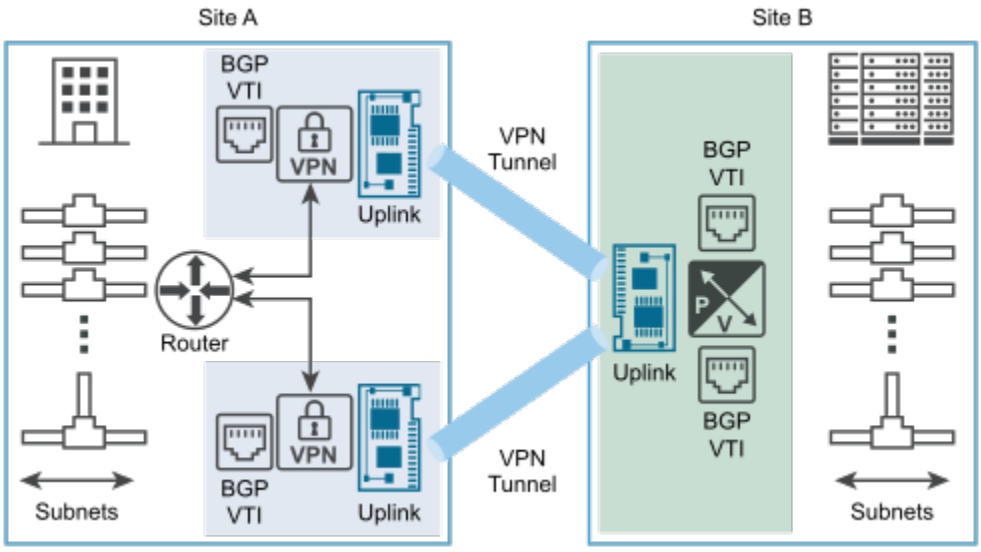
You can configure VPN tunnel redundancy with a route-based IPSec VPN session that is configured on a Tier-0 gateway. With tunnel redundancy, multiple tunnels can be set up between two sites, with one tunnel being used as the primary with failover to the other tunnels when the primary tunnel becomes unavailable. This feature is most useful when a site has multiple connectivity options, such as with different ISPs for link redundancy.

Important

- In NSX, IPSec VPN tunnel redundancy is supported using BGP only.
- Do not use static routing for route-based IPSec VPN tunnels to achieve VPN tunnel redundancy.

The following figure shows a logical representation of IPSec VPN tunnel redundancy between two sites. In this figure, Site A and Site B represent two data centers. For this example, assume that NSX is not managing the Edge VPN Gateways in Site A, and that NSX is managing an Edge Gateway virtual appliance in Site B.

Figure 7-1. Tunnel Redundancy in Route-Based IPSec VPN



As shown in the figure, you can configure two independent IPSec VPN tunnels by using VTIs. Dynamic routing is configured using BGP protocol to achieve tunnel redundancy. If both IPSec VPN tunnels are available, they remain in service. All the traffic destined from Site A to Site B through the NSX Edge node is routed through the VTI. The data traffic undergoes IPSec processing and goes out of its associated NSX Edge node uplink interface. All the incoming IPSec traffic received from Site B VPN Gateway on the NSX Edge node uplink interface is forwarded to the VTI after decryption, and then usual routing takes place.

You must configure BGP HoldDown timer and KeepAlive timer values to detect loss of connectivity with peer within the required failover time. See Configure BGP.

Understanding Layer 2 VPN

With Layer 2 VPN (L2 VPN), you can extend Layer 2 networks (VNIs or VLANs) across multiple sites on the same broadcast domain. This connection is secured with a route-based IPSec tunnel between the L2 VPN server and the L2 VPN client.

Note This L2 VPN feature is available only for NSX and does not have any third-party interoperability.

The extended network is a single subnet with a single broadcast domain, which means the VMs remain on the same subnet when they are moved between sites. The VMs' IP addresses do not change when they are moved. So, enterprises can seamlessly migrate VMs between network sites. The VMs can run on either VNI-based networks or VLAN-based networks. For cloud providers, L2 VPN provides a mechanism to onboard tenants without modifying existing IP addresses used by their workloads and applications.

In addition to supporting data center migration, an on-premises network extended with an L2 VPN is useful for a disaster recovery plan and dynamically engaging off-premise compute resources to meet the increased demand.

L2 VPN services are supported on both Tier-0 and Tier-1 gateways. Only one L2 VPN service (either client or server) can be configured for either Tier-0 or Tier-1 gateway.

Each L2 VPN session has one Generic Routing Encapsulation (GRE) tunnel. Tunnel redundancy is not supported. An L2 VPN session can extend up to 4094 L2 segments.

VLAN-based and VNI-based segments can be extended using L2 VPN service on an NSX Edge node that is managed in an NSX environment. You can extend L2 networks from VLAN to VNI, VLAN to VLAN, and VNI to VNI.

Segments can be connected to either Tier-0 or Tier-1 gateways and use L2 VPN services.

Also supported is VLAN trunking using virtual distributed switching (VDS) 7.0 or later running NSX. If there are sufficient compute and I/O resources, an NSX Edge cluster can extend multiple VLAN networks over a single interface using VLAN trunking.

Beginning with NSX 3.0, the L2 VPN path MTU discovery (PMTUD) feature is enabled by default. With the PMTUD enabled, the source host learns the path MTU value for the destination host through the L2 VPN tunnel and limits the length of the outgoing IP packet to the learned value. This feature helps avoid IP fragmentation and reassembly within the tunnel, as a result improving the L2 VPN performance.

The L2 VPN PMTUD feature is not applicable for non-IP, non-unicast, and unicast packets with the DF (Don't Fragment) flag cleared. The global PMTU cache timer expires every 10 minutes. To disable or enable L2 VPN PMTUD feature, see Enable and Disable L2 VPN Path MTU Discovery.

The L2 VPN service support is provided in the following deployment scenarios.

- Between an NSX L2 VPN server and an L2 VPN client hosted on an NSX Edge that is managed in an NSX Data Center for vSphere environment. A managed L2 VPN client supports both VLANs and VNIs.
- Between an NSX L2 VPN server and an L2 VPN client hosted on a standalone or unmanaged NSX Edge. An unmanaged L2 VPN client supports VLANs only.
- Between an NSX L2 VPN server and an L2 VPN client hosted on an autonomous NSX Edge. An autonomous L2 VPN client supports VLANs only.
- Beginning with NSX 2.4 release, L2 VPN service support is available between an NSX L2 VPN server and NSX L2 VPN clients. In this scenario, you can extend the logical L2 segments between two on-premises software-defined data centers (SDDCs).

The following table lists the compatible NSX versions that can be used for the L2 VPN server and client.

Table 7-2. NSX L2 VPN Client

L2 VPN Server Version (NSX)	L2 VPN Client Version (NSX) Validated	L2 VPN Client Version (NSX) Supported Not Validated
4.0.1.1	4.0.1.1, 4.0.0.1, 3.2.1.2	N/A
3.2.0	3.2.0, 3.1.3, 3.1.2	3.1.x and later
3.1.3	3.1.3, 3.1.2, 3.1.1	3.0.x and later
3.1.2	3.1.2, 3.1.1, 2.5.3	3.0.x and later
3.1.1	3.1.1, 3.1.0, 3.0.1	3.0.x and later
3.1.0	3.1.0, 3.0.1, 3.0.0	3.0.x and later
3.0.3	3.0.3, 3.0.2, 3.0.1	2.5.x and later
3.0.2	3.0.2, 3.0.1, 2.5.2	2.5.x and later
3.0.0	3.0.0, 2.5.0, 2.5.1	2.5.x and later

The following table lists the compatible NSX and NSX-v versions that can be used for the L2 VPN server and client.

Table 7-3. NSX for vSphere L2VPN Client

L2 VPN Server Version (NSX)	L2 VPN Client Version (NSX-v) Validated	L2 VPN Client Version (NSX-v) Supported Not Validated
3.2.x	6.4.12	6.4.x and later
3.1.x	6.4.8	6.4.x and later

Enable and Disable L2 VPN Path MTU Discovery

You can enable or disable the L2 VPN path MTU (PMTU) discovery feature using CLI commands. By default L2 VPN PMTU discovery is enabled.

Prerequisites

You must have the user name and password for the admin account to log in to the NSX Edge node.

Procedure

1 Log in with admin privileges to the CLI of the NSX Edge node .

2 To check the status of the L2 VPN PMTU discovery feature, use the following command.

```
Nsxedge> get dataplane l2vpn-pmtu config
```

If the feature is enabled, you see the following output: `l2vpn_pmtu_enabled : True`.

If the feature is disabled, you see the following output: `l2vpn_pmtu_enabled : False`.

3 To disable the L2 VPN PMTU discovery feature, use the following command.

```
nsxedge> set dataplane l2vpn-pmtu disabled
```

Adding VPN Services

You can add either an IPSec VPN (policy-based or route-based) or an L2 VPN using the NSX Manager user interface (UI).

The following sections provide information about the workflows required to set up the VPN service that you need. The topics that follow these sections provide details on how to add either an IPSec VPN or an L2 VPN using the NSX Manager user interface.

Policy-Based IPSec VPN Configuration Workflow

Configuring a policy-based IPSec VPN service workflow requires the following high-level steps.

1 Create and enable an IPSec VPN service using an existing Tier-0 or Tier-1 gateway. See Add an IPSec VPN Service.

2 Create a DPD (dead peer detection) profile, if you prefer not to use the system default. See Add DPD Profiles.

3 To use a non-system default IKE profile, define an IKE (Internet Key Exchange) profile. See Add IKE Profiles.

4 Configure an IPSec profile using Add IPSec Profiles.

5 Use Add Local Endpoints to create a VPN server hosted on the NSX Edge.

6 Configure a policy-based IPSec VPN session, apply the profiles, and attach the local endpoint to it. See Add a Policy-Based IPSec Session. Specify the local and peer subnets to be used for the tunnel. Traffic from a local subnet destined to the peer subnet is protected using the tunnel defined in the session.

Route-Based IPSec VPN Configuration Workflow

A route-based IPSec VPN configuration workflow requires the following high-level steps.

1 Configure and enable an IPSec VPN service using an existing Tier-0 or Tier-1 gateway. See Add an IPSec VPN Service.

2 Define an IKE profile if you prefer not to use the default IKE profile. See Add IKE Profiles.

3 If you decide not to use the system default IPSec profile, create one using Add IPSec Profiles.

4 Create a DPD profile if you want to do not want to use the default DPD profile. See Add DPD Profiles.

5 Add a local endpoint using Add Local Endpoints.

6 Configure a route-based IPSec VPN session, apply the profiles, and attach the local endpoint to the session. Provide a VTI IP in the configuration and use the same IP to configure routing. The routes can be static or dynamic (using BGP). See Add a Route-Based IPSec Session.

L2 VPN Configuration Workflow

Configuring an L2 VPN requires that you configure an L2 VPN service in Server mode and then another L2 VPN service in Client mode. You also must configure the sessions for the L2 VPN server and L2 VPN client using the peer code generated by the L2 VPN Server. Following is a high-level workflow for configuring an L2 VPN service.

1 Create an L2 VPN Service in Server mode.

   a Configure a route-based IPSec VPN tunnel with a Tier-0 or Tier-1 gateway and an L2 VPN Server service using that route-based IPSec tunnel. See Add an L2 VPN Server Service.

   b Configure an L2 VPN server session, which binds the newly created route-based IPSec VPN service and the L2 VPN server service, and automatically allocates the GRE IP addresses. See Add an L2 VPN Server Session.

   c Add segments to the L2 VPN Server sessions. This step is also described in Add an L2 VPN Server Session.

   d Use Download the Remote Side L2 VPN Configuration File to obtain the peer code for the L2 VPN Server service session, which must be applied on the remote site and used to configure the L2 VPN Client session automatically.

2 Create an L2 VPN Service in Client mode.

   a Configure another route-based IPSec VPN service using a different Tier-0 or Tier-1 gateway and configure an L2 VPN Client service using that Tier-0 or Tier-1 gateway that you just configured. See Add an L2 VPN Client Service for information.

   b Define the L2 VPN Client sessions by importing the peer code generated by the L2 VPN Server service. See Add an L2 VPN Client Session.

   c Add segments to the L2 VPN Client sessions defined in the previous step. This step is described in Add an L2 VPN Client Session.

Add an IPSec VPN Service

NSX supports a site-to-site IPSec VPN service between a Tier-0 or Tier-1 gateway and remote sites. You can create a policy-based or a route-based IPSec VPN service. You must create the IPSec VPN service first before you can configure either a policy-based or a route-based IPSec VPN session.

Note IPSec VPN is not supported in the NSX limited export release.

IPSec VPN is not supported when the local endpoint IP address goes through NAT in the same logical router that the IPSec VPN session is configured.

Prerequisites

- Familiarize yourself with the IPSec VPN. See Understanding IPSec VPN.
- You must have at least one Tier-0 or Tier-1 gateway configured and available for use. See Add a Tier-0 Gateway or Add a Tier-1 Gateway for more information.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Navigate to **Networking > VPN > VPN Services**.

3 Select **Add Service > IPSec**.

4 Enter a name for the IPSec service.

This name is required.

5 From the **Tier-0/Tier-1 Gateway** drop-down menu, select the Tier-0 or Tier-1 gateway to associate with this IPSec VPN service.

6 Enable or disable **Admin Status**.

By default, the value is set to `Enabled`, which means the IPSec VPN service is enabled on the Tier-0 or Tier-1 gateway after the new IPSec VPN service is configured.

7 Set the value for **IKE Log Level**.

The default is set to the `Info` level.

8 Enter a value for **Tags** if you want to include this service in a tag group.

9 To enable or disable the stateful synchronization of VPN sessions, toggle **Session sync**.

By default, the value is set to `Enabled`.

10 Click **Global Bypass Rules** if you want to allow data packets to be exchanged between the specified local and remote IP addresses without any IPSec protection. In the **Local Networks** and **Remote Networks** text boxes, enter the list of local and remote subnets between which the bypass rules are applied.

If you enable these rules, data packets are exchanged between the specified local and remote IP sites even if their IP addresses are specified in the IPSec session rules. The default is to use the IPSec protection when data is exchanged between local and remote sites. These rules apply for all IPSec VPN sessions created within this IPSec VPN service.

11 Click **Save**.

After the new IPSec VPN service is created successfully, you are asked whether you want to continue with the rest of the IPSec VPN configuration. If you click **Yes**, you are taken back to the Add IPSec VPN Service panel. The **Sessions** link is now enabled and you can click it to add an IPSec VPN session.

Results

After one or more IPSec VPN sessions are added, the number of sessions for each VPN service will appear in the **VPN Services** tab. You can reconfigure or add sessions by clicking the number in the **Sessions** column. You do not need to edit the service. If the number is zero, it is not clickable and you must edit the service to add sessions.

What to do next

Use information in Adding IPSec VPN Sessions to guide you in adding an IPSec VPN session. You also provide information for the profiles and local endpoint that are required to finish the IPSec VPN configuration.

Add an L2 VPN Service

You configure an L2 VPN service on a Tier-0 or Tier-1 gateway. To enable the L2 VPN service, you must first create an IPSec VPN service on the Tier-0 or Tier-1 gateway, if it does not exist yet. You then configure an L2 VPN tunnel between an L2 VPN server (destination gateway) and an L2 VPN client (source gateway).

To configure an L2 VPN service, use the information in the topics that follow in this section.

Prerequisites

- Familiarize yourself with IPsec VPN and L2 VPN. See Understanding IPSec VPN and Understanding Layer 2 VPN.
- You must have at least one Tier-0 or Tier-1 gateway configured and available for use. See Add a Tier-0 Gateway or Add a Tier-1 Gateway.

Procedure

1 Add an L2 VPN Server Service

   To configure an L2 VPN Server service, you must configure the L2 VPN service in server mode on the destination NSX Edge to which the L2 VPN client is to be connected.

2 Add an L2 VPN Client Service

   After configuring the L2 VPN Server service, configure the L2 VPN service in the client mode on another NSX Edge instance.

Add an L2 VPN Server Service

To configure an L2 VPN Server service, you must configure the L2 VPN service in server mode on the destination NSX Edge to which the L2 VPN client is to be connected.

Procedure

1 With admin privileges, log in to NSX Manager.

2 (Optional) If an IPSec VPN service does not exist yet on either a Tier-0 or Tier-1 gateway that you want to configure as the L2 VPN server, create the service using the following steps.

   a Navigate to the **Networking > VPN > VPN Services** tab and select **Add Service > IPSec**.

   b Enter a name for the IPSec VPN service.

   c From the **Tier-0/Tier-1 Gateway** drop-down menu, select the gateway to use with the L2 VPN server.

   d If you want to use values different from the system defaults, set the rest of the properties on the Add IPSec Service pane, as needed.

   e Click **Save** and when prompted if you want to continue configuring the IPSec VPN service, select **No**.

3 Navigate to the **Networking > VPN > VPN Services** tab and select **Add Service > L2 VPN Server** to create an L2 VPN server.

4 Enter a name for the L2 VPN server.

5 From the **Tier-0/Tier-1 Gateway** drop-down menu, select the same Tier-0 or Tier-1 gateway that you used with the IPSec service you created a moment ago.

6 Enter an optional description for this L2 VPN server.

7 Enter a value for **Tags** if you want to include this service in a tag group.

8 Enable or disable the **Hub & Spoke** property.

   By default, the value is set to `Disabled`, which means the traffic received from the L2 VPN clients is only replicated to the segments connected to the L2 VPN server. If this property is set to `Enabled`, the traffic from any L2 VPN client is replicated to all other L2 VPN clients.

9 Click **Save**.

   After the new L2 VPN server is created successfully, you are asked whether you want to continue with the rest of the L2 VPN service configuration. If you click **Yes**, you are taken back to the Add L2 VPN Server pane and the **Session** link is enabled. You can use that link to create an L2 VPN server session or use the **Networking > VPN > L2 VPN Sessions** tab.

Results

After one or more L2 VPN sessions are added, the number of sessions for each VPN service will appear in the **VPN Services** tab. You can reconfigure or add sessions by clicking the number in the **Sessions** column. You do not need to edit the service. Note that if the number is zero, it is not clickable and you must edit the service to add sessions.

What to do next

Configure an L2 VPN server session for the L2 VPN server that you configured using information in Add an L2 VPN Server Session as a guide.

Add an L2 VPN Client Service

After configuring the L2 VPN Server service, configure the L2 VPN service in the client mode on another NSX Edge instance.

Procedure

1 With admin privileges, log in to NSX Manager.

2 (Optional) If an IPSec VPN service does not exist yet on either a Tier-0 or Tier-1 gateway that you want to configure as the L2 VPN client, create the service using the following steps.

   a Navigate to the **Networking > VPN > VPN Services** tab and select **Add Service > IPSec**.

   b Enter a name for the IPSec VPN service.

   c From the **Tier-0/Tier-1 Gateway** drop-down menu, select a Tier-0 or Tier-1 gateway to use with the L2 VPN client.

   d If you want to use values different from the system defaults, set the rest of the properties on the Add IPSec Service pane, as needed.

   e Click **Save** and when prompted if you want to continue configuring the IPSec VPN service, select **No**.

3 Navigate to the **Networking > VPN > VPN Services** tab and select **Add Service > L2 VPN Client**.

4 Enter a name for the L2 VPN Client service.

5 From the **Tier-0/Tier-1 Gateway** drop-down menu, select the same Tier-0 or Tier-1 gateway that you used with the route-based IPSec tunnel you created a moment ago.

6 Optionally set the values for **Description** and **Tags**.

7 Click **Save**.

   After the new L2 VPN client service is created successfully, you are asked whether you want to continue with the rest of the L2 VPN client configuration. If you click **Yes**, you are taken back to the Add L2 VPN Client pane and the **Session** link is enabled. You can use that link to create an L2 VPN client session or use the **Networking > VPN > L2 VPN Sessions** tab.

Results

After one or more L2 VPN sessions are added, the number of sessions for each VPN service will appear in the **VPN Services** tab. You can reconfigure or add sessions by clicking the number in the **Sessions** column. You do not need to edit the service. If the number is zero, it is not clickable and you must edit the service to add sessions.

What to do next

Configure an L2 VPN client session for the L2 VPN Client service that you configured. Use the information in Add an L2 VPN Client Session as a guide.

Adding IPSec VPN Sessions

After you have configured an IPSec VPN service, you must add either a policy-based IPSec VPN session or a route-based IPSec VPN session, depending on the type of IPSec VPN you want to configure. You also provide the information for the local endpoint and profiles to use to finish the IPSec VPN service configuration.

Using Certificate-Based Authentication for IPSec VPN Sessions

When you use certificate-based authentication for an IPSec VPN session, you must configure the certificate details for the IPSec session in the associated local endpoint.

Note Wildcard certificates are not supported for IPSec VPN.

Refer to the following workflow for details on how to configure the certificate details for a IPSec VPN session.

Configure Certificate-Based Authentication for an IPSec VPN Session

1 Create and enable an IPSec VPN service using an existing Tier-0 or Tier-1 gateway. See Add an IPSec VPN Service.

2 If you do not have the necessary server certificates or CA certificates in NSX Manager, import the certificates. See Import a Self-signed or CA-signed Certificate and Import a CA Certificate.

3 Use Add Local Endpoints to create a VPN server hosted on the logical router and select the certificates for it.

The local ID is derived from the certificate associated with the local endpoint and depends on the X509v3 extensions present in the certificate. The local ID can be either the X509v3 extension Subject Alternative Name (SAN) or Distinguished Name (DN). The **Local ID** is not required and the ID specified there is ignored. However, for the remote VPN gateway, you need to configure the local ID as remote ID in the peer VPN gateway.

- If `X509v3 Subject Alternative Name` is found in the certificate, then one of the SAN strings is taken as the local ID value.

  If the certificate has multiple SAN fields, then following order is used to select the local ID.

Order	SAN Field
1	IP Address
2	DNS
3	Email Address

For example, if the configured site certificate has the following SAN fields,

```
X509v3 Subject Alternative Name:
DNS:Site123.vmware.com, email:user1@company.com, IP Address:1.1.1.1
```

then the IP address `1.1.1.1` is used as the local ID. If the IP address is not available, then the DNS string is used. And if the IP address and the DNS are not available, then the email address is used.

- If `X509v3 Subject Alternative Name` is not present in the certificate, then the Distinguished Name (DN) is used as the local ID value.

  For example, if the certificate does not have any SAN fields, and its DN string is

  ```
  C=US, ST=California, O=MyCompany, OU=MyOrg, CN=Site123
  ```

  then the DN string automatically becomes the local ID. The local ID is the peer ID on the remote site.

Note If the certificate details are not properly configured, it might cause the VPN session to go down with the `Down` alarm of `Authentication failed`.

4 Configure either a policy-based or route-based IPSec VPN session. See Add a Policy-Based IPSec Session or Add a Route-Based IPSec Session.

Make sure to configure the following settings.

a From the **Authentication Mode** drop-down menu, select **Certificate**.

b In the **Remote ID** textbox, enter a value to identify the peer site.

The remote ID must be a distinguished name (DN), IP address, DNS, or an email address used in the peer site's certificate.

Note If the peer site's certificate contains an email address in the DN string, for example,

```
C=US, ST=California, O=MyCompany, OU=MyOrg, CN=Site123/emailAddress=user1@mycompany.com
```

then enter the **Remote ID** value using the following format as an example.

```
C=US, ST=California, O=MyCompany, OU=MyOrg, CN=Site123, MAILTO=user1@mycompany.com
```

Add a Policy-Based IPSec Session

When you add a policy-based IPSec VPN, IPSec tunnels are used to connect multiple local subnets that are behind the NSX Edge node with peer subnets on the remote VPN site.

The following steps use the **IPSec Sessions** tab on the NSX Manager UI to create a policy-based IPSec session. You also add information for the tunnel, IKE, and DPD profiles, and select an existing local endpoint to use with the policy-based IPSec VPN.

Note You can also add the IPSec VPN sessions immediately after you have successfully configured the IPSec VPN service. You click **Yes** when prompted to continue with the IPSec VPN service configuration and select **Sessions > Add Sessions** on the Add IPsec Service panel. The first few steps in the following procedure assume you selected **No** to the prompt to continue with the IPSec VPN service configuration. If you selected **Yes**, proceed to step 3 in the following steps to guide you with the rest of the policy-based IPSec VPN session configuration.

Prerequisites

- You must have configured an IPSec VPN service before proceeding. See Add an IPSec VPN Service.
- Obtain the information for the local endpoint, IP address for the peer site, local network subnet, and remote network subnet to use with the policy-based IPSec VPN session you are adding. To create a local endpoint, see Add Local Endpoints.
- If you are using a Pre-Shared Key (PSK) for authentication, obtain the PSK value.
- If you are using a certificate for authentication, ensure that the necessary server certificates and corresponding CA-signed certificates are already imported. See Chapter 23 Certificates.
- If you do not want to use the defaults for the IPSec tunnel, IKE, or dead peer detection (DPD) profiles provided by NSX, configure the profiles you want to use instead. See Adding Profiles for information.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Navigate to the **Networking > VPN > IPSec Sessions** tab.
3. Select **Add IPSec Session > Policy Based**.
4. Enter a name for the policy-based IPSec VPN session.
5. From the **VPN Service** drop-down menu, select the IPSec VPN service to which you want to add this new IPSec session.

   Note If you are adding this IPSec session from the **Add IPSec Sessions** dialog box, the VPN Service name is already indicated above the **Add IPSec Session** button.

6. Select an existing local endpoint from the drop-down menu.

   This local endpoint value is required and identifies the local NSX Edge node. If you want to create a different local endpoint, click the three-dot menu () and select **Add Local Endpoint**.

7 In the **Remote IP** text box, enter the required IP address of the remote site.

This value is required.

8 Enter an optional description for this policy-based IPSec VPN session.

The maximum length is 1024 characters.

9 To enable or disable the IPSec VPN session, click **Admin Status**.

By default, the value is set to `Enabled`, which means the IPSec VPN session is to be configured down to the NSX Edge node.

10 (Optional) From the **Compliance suite** drop-down menu, select a security compliance suite.

Note Compliance suite support is provided beginning with NSX 2.5. See About Supported Compliance Suites for more information.

The default value selected is `None`. If you select a compliance suite, the **Authentication Mode** is set to `Certificate` and in the **Advanced Properties** section, the values for **IKE profile** and **IPSec profile** are set to the system-defined profiles for the selected security compliance suite. You cannot edit these system-defined profiles.

11 If the **Compliance Suite** is set to `None`, select a mode from the **Authentication Mode** drop-down menu.

The default authentication mode used is `PSK`, which means a secret key shared between NSX Edge and the remote site is used for the IPSec VPN session. If you select `Certificate`, the site certificate that was used to configure the local endpoint is used for authentication.

For more information about certificate-based authentication, see Using Certificate-Based Authentication for IPSec VPN Sessions.

12 In the Local Networks and Remote Networks text boxes, enter at least one IP subnet address to use for this policy-based IPSec VPN session.

These subnets must be in a CIDR format.

13 If **Authentication Mode** is set to `PSK`, enter the key value in the **Pre-shared Key** text box.

This secret key can be a string with a maximum length of 128 characters.

Caution Be careful when sharing and storing a PSK value because it contains some sensitive information.

14 To identify the peer site, enter a value in **Remote ID**.

For peer sites using PSK authentication, this ID value must be the public IP address or the FQDN of the peer site. For peer sites using certificate authentication, this ID value must be the common name (CN) or distinguished name (DN) used in the peer site's certificate.

Note If the peer site's certificate contains an email address in the DN string, for example,

```
C=US, ST=California, O=MyCompany, OU=MyOrg, CN=Site123/emailAddress=user1@mycompany.com
```

then enter the **Remote ID** value using the following format as an example.

```
C=US, ST=California, O=MyCompany, OU=MyOrg, CN=Site123, MAILTO=user1@mycompany.com"
```

If the local site's certificate contains an email address in the DN string and the peer site uses the strongSwan IPsec implementation, enter the local site's ID value in that peer site. The following is an example.

```
C=US, ST=California, O=MyCompany, OU=MyOrg, CN=Site123, E=user1@mycompany.com"
```

15 To change the profiles, initiation mode, TCP MSS clamping mode, and tags used by the policy-based IPSec VPN session, click **Advanced Properties**.

By default, the system generated profiles are used. Select another available profile if you do not want to use the default. If you want to use a profile that is not configured yet, click the three-dot menu () to create another profile. See Adding Profiles.

a If the **IKE Profiles** drop-down menu is enabled, select the IKE profile.

b Select the IPsec tunnel profile, if the **IPSec Profiles** drop-down menu is not disabled.

c Select the preferred DPD profile if the **DPD Profiles** drop-down menu is enabled.

d Select the preferred mode from the **Connection Initiation Mode** drop-down menu.

Connection initiation mode defines the policy used by the local endpoint in the process of tunnel creation. The default value is **`Initiator`**. The following table describes the different connection initiation modes available.

Table 7-4. Connection Initiation Modes

Connection Initiation Mode	Description
`Initiator`	The default value. In this mode, the local endpoint initiates the IPSec VPN tunnel creation and responds to incoming tunnel setup requests from the peer gateway.
`On Demand`	In this mode, the local endpoint initiates the IPSec VPN tunnel creation after the first packet matching the policy rule is received. It also responds to the incoming initiation request.
`Respond Only`	The IPSec VPN never initiates a connection. The peer site always initiates the connection request and the local endpoint responds to that connection request.

e If you want to reduce the maximum segment size (MSS) payload of the TCP session during the IPSec connection, enable **TCP MSS Clamping**, select the **TCP MSS direction** value, and optionally set the **TCP MSS Value**.

See Understanding TCP MSS Clamping for more information.

f If you want to include this session as part of a specific group, enter the tag name in **Tags**.

16 Click **Save**.

Results

When the new policy-based IPSec VPN session is configured successfully, it is added to the list of available IPsec VPN sessions. It is in read-only mode.

What to do next

- Verify that the IPSec VPN tunnel status is Up. See Monitor and Troubleshoot VPN Sessions for information.
- If necessary, manage the IPSec VPN session information by clicking the three-dot menu () on the left-side of the session's row. Select one of the actions you are allowed to perform.

Add a Route-Based IPSec Session

When you add a route-based IPSec VPN, tunneling is provided on traffic that is based on routes that were learned dynamically over a virtual tunnel interface (VTI) using a preferred protocol, such as BGP. IPSec secures all the traffic flowing through the VTI.

The steps described in this topic use the **IPSec Sessions** tab to create a route-based IPSec session. You also add information for the tunnel, IKE, and DPD profiles, and select an existing local endpoint to use with the route-based IPSec VPN.

Note You can also add the IPSec VPN sessions immediately after the IPSec VPN service is successfully configured. Click **Yes** when prompted to continue with the IPSec VPN service configuration and select **Sessions > Add Sessions** on the Add IPSec Service panel. The first few steps in the following procedure assume you selected **No** to the prompt to continue with the IPSec VPN service configuration. If you selected **Yes**, proceed to step 3 to guide you with the rest of the route-based IPSec VPN session configuration.

Prerequisites

- You must have configured an IPSec VPN service before proceeding. See Add an IPSec VPN Service.
- Obtain the information for the local endpoint, IP address for the peer site, and tunnel service IP subnet address to use with the route-based IPSec session you are adding. To create a local endpoint, see Add Local Endpoints.
- If you are using a Pre-Shared Key (PSK) for authentication, obtain the PSK value.
- If you are using a certificate for authentication, ensure that the necessary server certificates and corresponding CA-signed certificates are already imported. See Chapter 23 Certificates.
- If you do not want to use the default values for the IPSec tunnel, IKE, or dead peer detection (DPD) profiles provided by NSX, configure the profiles you want to use instead. See Adding Profiles for information.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Navigate to **Networking > VPN > IPSec Sessions**.
3. Select **Add IPSec Session > Route Based**.
4. Enter a name for the route-based IPSec session.
5. From the **VPN Service** drop-down menu, select the IPSec VPN service to which you want to add this new IPSec session.

   Note If you are adding this IPSec session from the **Add IPSec Sessions** dialog box, the VPN Service name is already indicated above the **Add IPSec Session** button.

6. Select an existing local endpoint from the drop-down menu.

   The local endpoint value is required and identifies the local NSX Edge node. If you want to create a different local endpoint, click the three-dot menu () and select **Add Local Endpoint**.

7 In the **Remote IP** text box, enter the IP address of the remote site.

This is a required value.

8 Enter an optional description for this route-based IPSec VPN session.

The maximum length is 1024 characters.

9 To enable or disable the IPSec session, click **Admin Status**.

By default, the value is set to `Enabled`, which means the IPSec session is to be configured down to the NSX Edge node.

10 (Optional) From the **Compliance suite** drop-down menu, select a security compliance suite.

Note Compliance suite support is provided beginning with NSX 2.5. See About Supported Compliance Suites for more information.

The default value is set to `None`. If you select a compliance suite, the **Authentication Mode** is set to `Certificate` and in the **Advanced Properties** section, the values for **IKE profile** and **IPSec profile** are set to the system-defined profiles for the selected compliance suite. You cannot edit these system-defined profiles.

11 Enter an IP subnet address in **Tunnel Interface** in the CIDR notation.

This address is required.

12 If the **Compliance Suite** is set to `None`, select a mode from the **Authentication Mode** drop-down menu.

The default authentication mode used is `PSK`, which means a secret key shared between NSX Edge and the remote site is used for the IPSec VPN session. If you select `Certificate`, the site certificate that was used to configure the local endpoint is used for authentication.

For more information about certificate-based authentication, see Using Certificate-Based Authentication for IPSec VPN Sessions.

13 If you selected `PSK` for the authentication mode, enter the key value in the **Pre-shared Key** text box.

This secret key can be a string with a maximum length of 128 characters.

Caution Be careful when sharing and storing a PSK value because it contains some sensitive information.

14 Enter a value in **Remote ID**.

For peer sites using PSK authentication, this ID value must be the public IP address or the FQDN of the peer site. For peer sites using certificate authentication, this ID value must be the common name (CN) or distinguished name (DN) used in the peer site's certificate.

Note If the peer site's certificate contains an email address in the DN string, for example,

```
C=US, ST=California, O=MyCompany, OU=MyOrg, CN=Site123/emailAddress=user1@mycompany.com
```

then enter the **Remote ID** value using the following format as an example.

```
C=US, ST=California, O=MyCompany, OU=MyOrg, CN=Site123, MAILTO=user1@mycompany.com"
```

If the local site's certificate contains an email address in the DN string and the peer site uses the strongSwan IPsec implementation, enter the local site's ID value in that peer site. The following is an example.

```
C=US, ST=California, O=MyCompany, OU=MyOrg, CN=Site123, E=user1@mycompany.com"
```

15 If you want to include this IPSec session as part of a specific group tag, enter the tag name in **Tags**.

16 To change the profiles, initiation mode, TCP MSS clamping mode, and tags used by the route-based IPSec VPN session, click **Advanced Properties**.

By default, the system-generated profiles are used. Select another available profile if you do not want to use the default. If you want to use a profile that is not configured yet, click the three-dot menu () to create another profile. See Adding Profiles.

a If the **IKE Profiles** drop-down menu is enabled, select the IKE profile.

b Select the IPsec tunnel profile, if the **IPSec Profiles** drop-down menu is not disabled.

c Select the preferred DPD profile if the **DPD Profiles** drop-down menu is enabled.

d Select the preferred mode from the **Connection Initiation Mode** drop-down menu.

Connection initiation mode defines the policy used by the local endpoint in the process of tunnel creation. The default value is **`Initiator`**. The following table describes the different available connection initiation modes.

Table 7-5. Connection Initiation Modes

Connection Initiation Mode	Description
`Initiator`	The default value. In this mode, the local endpoint initiates the IPSec VPN tunnel creation and responds to incoming tunnel setup requests from the peer gateway.
`On Demand`	Do not use with the route-based VPN. This mode applies to policy-based VPN only.
`Respond Only`	The IPSec VPN never initiates a connection. The peer site always initiates the connection request and the local endpoint responds to that connection request.

17 If you want to reduce the maximum segment size (MSS) payload of the TCP session during the IPSec connection, enable **TCP MSS Clamping**, select the **TCP MSS** direction value, and optionally set the **TCP MSS Value**.

See Understanding TCP MSS Clamping for more information.

18 If you want to include this IPSec session as part of a specific group tag, enter the tag name in **Tags**.

19 Click **Save**.

Results

When the new route-based IPSec VPN session is configured successfully, it is added to the list of available IPsec VPN sessions. It is in read-only mode.

What to do next

- Verify that the IPSec VPN tunnel status is Up. See Monitor and Troubleshoot VPN Sessions for information.
- Configure routing using either a static route or BGP. See Configure a Static Route or Configure BGP.
- If necessary, manage the IPSec VPN session information by clicking the three-dot menu ( ) on the left-side of the session's row. Select one of the actions you can perform.

About Supported Compliance Suites

You can specify a security compliance suite to use to configure the security profiles used for an IPSec VPN session.

A security compliance suite has predefined values that are used for different security parameters and that cannot be modified. When you select a compliance suite, the predefined values are automatically used for the security profile of the IPSec VPN session you are configuring.

The following table lists the compliance suites that are supported for IKE profiles in NSX and the values that are predefined for each.

Compliance Suite Name	IKE Version	Encryption Algorithm	Digest Algorithm	Diffie Hellman Group
CNSA	IKE V2	AES 256	SHA2 384	Group 15, Group 20
FIPS	IKE FLEX	AES 128	SHA2 256	Group 20
Foundation	IKE V1	AES 128	SHA2 256	Group 14
PRIME	IKE V2	AES GCM 128	Not Set	Group 19
Suite-B-GCM-128	IKE V2	AES 128	SHA2 256	Group 19
Suite-B-GCM-256	IKE V2	AES 256	SHA2 384	Group 20

Note The AES 128 and AES 256 algorithms use the CBC mode of operation.

The following table lists the compliance suites that are supported for IPSec profiles in NSX and the values that are predefined for each.

Compliance Suite Name	Encryption Algorithm	Digest Algorithm	PFS Group	Diffie-Hellman Group
CNSA	AES 256	SHA2 384	Enabled	Group 15, Group 20
FIPS	AES GCM 128	Not Set	Enabled	Group 20
Foundation	AES 128	SHA2 256	Enabled	Group 14
PRIME	AES GCM 128	Not Set	Enabled	Group 19
Suite-B-GCM-128	AES GCM 128	Not Set	Enabled	Group 19
Suite-B-GCM-256	AES GCM 256	Not Set	Enabled	Group 20

Note The AES 128 and AES 256 algorithms use the CBC mode of operation.

Adding L2 VPN Sessions

After you have configured an L2 VPN server and an L2 VPN client, you must add L2 VPN sessions for both to complete the L2 VPN service configuration.

Add an L2 VPN Server Session

After creating an L2 VPN Server service, you must add an L2 VPN session and attach it to an existing segment.

The following steps use the **L2 VPN Sessions** tab on the NSX Manager UI to create an L2 VPN Server session. You also select an existing local endpoint and segment to attach to the L2 VPN Server session.

Note You can also add an L2 VPN Server session immediately after you have successfully configured the L2 VPN Server service. You click **Yes** when prompted to continue with the L2 VPN Server configuration and select **Sessions > Add Sessions** on the Add L2 VPN Server panel. The first few steps in the following procedure assume you selected **No** to the prompt to continue with the L2 VPN Server configuration. If you selected **Yes**, proceed to step 3 in the following steps to guide you with the rest of the L2 VPN Server session configuration.

Prerequisites

- You must have configured an L2 VPN Server service before proceeding. See Add an L2 VPN Server Service.
- Obtain the information for the local endpoint and remote IP to use with the L2 VPN Server session you are adding. To create a local endpoint, see Add Local Endpoints.
- Obtain the values for the pre-shared key (PSK) and the tunnel interface subnet to use with the L2 VPN Server session.
- Obtain the name of the existing segment you want to attach to the L2 VPN Server session you are creating. See Add a Segment for information.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Navigate to the **Networking > VPN > L2 VPN Sessions** tab.
3. Select **Add L2 VPN Session > L2 VPN Server**.
4. Enter a name for the L2 VPN Server session.
5. From the **VPN Service** drop-down menu, select the IPsec service on the same Tier-0 gateway for which the L2 VPN session is being created.

   Note If you are adding this L2 VPN Server session from the Set L2VPN Server Sessions dialog box, the L2 VPN Server service is already indicated above the **Add L2 Session** button.
6. Select an existing local endpoint from the drop-down menu.

   If you want to create a different local endpoint, click the three-dot menu () and select **Add Local Endpoint**.
7. Enter the IP address of the remote site under Remote IP.
8. To enable or disable the L2 VPN Server session, click **Admin Status**.

   By default, the value is set to **`Enabled`**, which means the L2 VPN Server session is to be configured down to the NSX Edge node.

9 Enter the secret key value in **Pre-shared Key**.

Caution Be careful when sharing and storing a PSK value because it is considered sensitive information.

10 Enter an IP subnet address in the **Tunnel Interface** using the CIDR notation.

For example, 4.5.6.6/24. This subnet address is required.

11 Enter a value in **Remote ID**.

For peer sites using certificate authentication, this ID must be the common name in the peer site's certificate. For PSK peers, this ID can be any string. Preferably, use the public IP address of the VPN or an FQDN for the VPN services as the `Remote ID`.

12 If you want to include this session as part of a specific group, enter the tag name in **Tags**.

13 Click **Advanced Properties.** if you want to reduce the maximum segment size (MSS) payload of the TCP session during the L2 VPN connection.

By default, **TCP MSS Clamping** is enabled and the **TCP MSS Direction** is set to **`Both`**. See Understanding TCP MSS Clamping for more information.

a Enable or disable **TCP MSS Clamping**.

b Set the **TCP MSS Value**, if necessary. If the field is left blank, the value is automatically assigned.

14 Click **Save** and click **Yes** when prompted if you want to continue with the VPN service configuration.

You are returned to the Add L2VPN Sessions panel and the **Segments** link is now enabled.

15 Attach an existing segment to the L2 VPN Server session.

a Click **Segments > Set Segments**.

b In the **Set Segments** dialog box, click **Set Segment** to attach an existing segment to the L2 VPN Server session.

c From the **Segment** drop-down menu, select the VNI-based or VLAN-based segment that you want to attach to the session.

d Enter a unique value in the **VPN Tunnel ID** that is used to identify the segment that you selected.

e In the **Local Egress Gateway IP** text box, enter the IP address of the local gateway that your workload VMs on the segment use as their default gateway. The same IP address can be configured in the remote site on the extended segment.

f Click **Save** and then **Close**.

In the Set L2VPN Sessions pane or dialog box, the system has incremented the **Segments** count for the L2 VPN Server session.

16 To finish the L2 VPN Server session configuration, click **Close Editing**.

Results

In the **VPN Services** tab, the system incremented the **Sessions** count for the L2 VPN Server service that you configured.

If you have attached one or more segments to the session, you see the number of segments for each session in the **L2 VPN Sessions** tab. You can reconfigure or add segments by clicking the number in the **Segments** column. You do not need to edit the session. If the number is zero, it is not clickable and you must edit the session to add segments.

What to do next

To complete the L2 VPN service configuration, you must also create an L2 VPN service in Client mode and an L2 VPN client session. See Add an L2 VPN Client Service and Add an L2 VPN Client Session.

Add an L2 VPN Client Session

You must add an L2 VPN Client session after creating an L2 VPN Client service, and attach it to an existing segment.

The following steps use the **L2 VPN Sessions** tab on the NSX Manager UI to create an L2 VPN Client session. You also select an existing local endpoint and segment to attach to the L2 VPN Client session.

Note You can also add an L2 VPN Client session immediately after you have successfully configured the L2 VPN Client service. Click **Yes** when prompted to continue with the L2 VPN Client configuration and select **Sessions > Add Sessions** on the Add L2 VPN Client panel. The first few steps in the following procedure assume you selected **No** to the prompt to continue with the L2 VPN Client configuration. If you selected **Yes**, proceed to step 3 in the following steps to guide you with the rest of the L2 VPN Client session configuration.

Prerequisites

- You must have configured an L2 VPN Client service before proceeding. See Add an L2 VPN Client Service.
- Obtain the IP addresses information for the local IP and remote IP to use with the L2 VPN Client session you are adding.
- Obtain the peer code that was generated during the L2 VPN server configuration. See Download the Remote Side L2 VPN Configuration File.
- Obtain the name of the existing segment you want to attach to the L2 VPN Client session you are creating. See Add a Segment.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Select the **Networking > VPN > L2 VPN Sessions**.

3 Select **Add L2 VPN Session > L2 VPN Client**.

4 Enter a name for the L2 VPN Client session.

5 From the **VPN Service** drop-down menu, select the L2 VPN Client service with which the L2 VPN session is to be associated.

Note If you are adding this L2 VPN Client session from the Set L2VPN Client Sessions dialog box, the L2 VPN Client service is already indicated above the **Add L2 Session** button.

6 In the **Local IP address** text box, enter the IP address of the L2 VPN Client session.

7 Enter the remote IP address of the IPSec tunnel to be used for the L2 VPN Client session.

8 In the **Peer Configuration** text box, enter the peer code generated when you configured the L2 VPN Server service.

9 Enable or disable **Admin Status**.

By default, the value is set to **`Enabled`**, which means the L2 VPN Server session is to be configured down to the NSX Edge node.

10 Click **Save** and click **Yes** when prompted if you want to continue with the VPN service configuration.

11 Attach an existing segment to the L2 VPN Client session.

a Select **Segments > Add Segments**.

b In the **Set Segments** dialog box, click **Add Segment**.

c From the **Segment** drop-down menu, select the VNI-based or VLAN-based segment you want to attach to the L2 VPN Client session.

d Enter a unique value in the **VPN Tunnel ID** that is used to identify the segment that you selected.

e Click **Close**.

12 To finish the L2 VPN Client session configuration, click **Close Editing**.

Results

In the **VPN Services** tab, the sessions count is updated for the L2 VPN Client service that you configured.

If you have attached one or more segments to the session, you see the number of segments for each session in the **L2 VPN Sessions** tab. You can reconfigure or add segments by clicking the number in the **Segments** column. You do not need to edit the session. If the number is zero, it is not clickable and you must edit the session to add segments.

Download the Remote Side L2 VPN Configuration File

To configure the L2 VPN client session, you must obtain the peer code that was generated when you configured the L2 VPN server session.

Prerequisites

- You must have configured an L2 VPN server service and a session successfully before proceeding. See Add an L2 VPN Server Service.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Navigate to the **Networking > VPN > L2 VPN Sessions** tab.

3 In the table of L2 VPN sessions, expand the row for the L2 VPN server session you plan to use for the L2 VPN client session configuration.

4 Click **Download Config** and click **Yes** on the Warning dialog box.

A text file with the name `L2VPNSession_<name-of-L2-VPN-server-session>_config.txt` is downloaded. It contains the peer code for the remote side L2 VPN configuration.

Caution Be careful when storing and sharing the peer code because it contains a PSK value, which is considered sensitive information.

For example, `L2VPNSession_L2VPNServer_config.txt` contains the following configuration.

```
[
  {
    "transport_tunnel_path": "/infra/tier-0s/ServerT0_AS/locale-services/1-
policyconnectivity-693/ipsec-vpn-services/IpsecService1/sessions/Routebase1",
    "peer_code":
"MCw3ZjBjYzdjLHsic2l0ZU5hbWUiOiJSb3V0ZWJhc2UxIiwic3JjVGFwSXAiOiIxNjkuMjU0LjY0LjIiLCJkc3RUYX
BJcCI6IjE2OS4yNTQuNjQuMSIsImlrZU9wdGl
vbiI6ImlrZXYyIiwiZW5jYXBQcm90byI6ImdyZS9pcHNlYyIsImRoR3JvdXAiOiJkaDE0IiwiZW5jcnlwdEFuZERpZ2
VzdCI6ImFlcy1nY20vc2hhLTI1NiIsInBzayI
6IlZNd2FyZTEyMyIsInR1bm5lbHMiOlt7ImxvY2FsSWQiOiI2MC42MC42MC4xIiwicGVlcklkIjoiNTAuNTAuNTAuMS
IsImxvY2FsVnRpSXAiOiIxNjkuMi4yLjMvMzEifV19"
  }
]
```

5 Copy the peer code, which you use to configure the L2 VPN client service and session.

Using the preceding configuration file example, the following peer code is what you copy to use with the L2 VPN client configuration.

```
MCw3ZjBjYzdjLHsic2l0ZU5hbWUiOiJSb3V0ZWJhc2UxIiwic3JjVGFwSXAiOiIxNjkuMjU0LjY0LjIiLCJkc3RUYXB
JcCI6IjE2OS4yNTQuNjQuMSIsImlrZU9wdGl
vbiI6ImlrZXYyIiwiZW5jYXBQcm90byI6ImdyZS9pcHNlYyIsImRoR3JvdXAiOiJkaDE0IiwiZW5jcnlwdEFuZERpZ2
VzdCI6ImFlcy1nY20vc2hhLTI1NiIsInBzayI
6IlZNd2FyZTEyMyIsInR1bm5lbHMiOlt7ImxvY2FsSWQiOiI2MC42MC42MC4xIiwicGVlcklkIjoiNTAuNTAuNTAuMS
IsImxvY2FsVnRpSXAiOiIxNjkuMi4yLjMvMzEifV19
```

What to do next

Configure the L2 VPN Client service and session. See Add an L2 VPN Client Service and Add an L2 VPN Client Session.

Add Local Endpoints

You must configure a local endpoint to use with the IPSec VPN that you are configuring.

The following steps use the **Local Endpoints** tab on the NSX Manager UI. You can also create a local endpoint while in the process of adding an IPSec VPN session by clicking the three-dot menu () and selecting **Add Local Endpoint**. If you are in the middle of configuring an IPSec VPN session, proceed to step 3 in the following steps to guide you with creating a new local endpoint.

Prerequisites

- If you are using a certificate-based authentication mode for the IPSec VPN session that is to use the local endpoint you are configuring, obtain the information about the certificate that the local endpoint must use.
- Ensure that you have configured an IPSec VPN service to which this local endpoint is to be associated.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Navigate to **Networking > VPN > Local Endpoints** and click **Add Local Endpoint**.
3. Enter a name for the local endpoint.
4. From the **VPN Service** drop-down menu, select the IPSec VPN service to which this local endpoint is to be associated.
5. Enter an IP address for the local endpoint.

   For an IPSec VPN service running on a Tier-0 gateway, the local endpoint IP address must be different from the Tier-0 gateway's uplink interface IP address. The local endpoint IP address you provide is associated with the loopback interface for the Tier-0 gateway and is also published as a routable IP address over the uplink interface. For IPSec VPN service running on a Tier-1 gateway, in order for the local endpoint IP address to be routable, the route advertisement for IPSec local endpoints must be enabled in the Tier-1 gateway configuration. See Add a Tier-1 Gateway for more information.
6. If you are using a certificate-based authentication mode for the IPSec VPN session, from the **Site Certificate** drop-down menu, select the certificate that is to be used by the local endpoint.
7. (Optional) Optionally add a description in **Description**.

8 Enter the **Local ID** value that is used for identifying the local NSX Edge instance.

This local ID is the peer ID on the remote site. The local ID must be either the public IP address or FQDN of the remote site. For IPSec VPN sessions with certificate-based authentication and are associated with the local endpoint, the **Local ID** is derived from the certificate associated with the local endpoint. The ID specified in the **Local ID** text box is ignored. The local ID derived from the certificate for a VPN session depends on the extensions present in the certificate.

- If the X509v3 extension `X509v3 Subject Alternative Name` is not present in the certificate, then the Distinguished Name (DN) is used as the local ID value.

  For example, if the certificate does not have any Subject Alternative Name (SAN) fields and its DN string is:

  ```
  C=US, ST=California, O=MyCompany, OU=MyOrg, CN=Site123
  ```

  then the DN string is used as the local ID. This local ID is the peer ID on the remote site.

- If the X509v3 extension `X509v3 Subject Alternative Name` is found in the certificate, then one of the SAN fields is taken as the local ID value.

  If the certificate has multiple SAN fields, then the following order is used to select the local ID.

Order	SAN Field
1	IP Address
2	DNS
3	Email Address

  For example, if the configured site certificate has the following SAN fields:

  ```
  x509v3 Subject Alternative Name:
  DNS:Site123.vmware.com, email:user@company.com, IP Address:1.1.1.1
  ```

  then the IP address `1.1.1.1` is used as the local ID. If the IP address is not available, then the DNS string is used. And if the IP address and DNS are not available, then the email address is used.

To see the local ID that is used for an IPSec VPN session, do the following:

a Navigate to **Networking > VPN** and then click the **IPSec Sessions** tab.

b Expand the IPSec VPN session.

c Click **Download Config** to download the configuration file which contains the local ID.

9 From the **Trusted CA Certificates** and **Certificate Revocation List** drop-down menus, select the appropriate certificates that are required for the local endpoint.

10 (Optional) Specify a tag.

11 Click **Save**.

Adding Profiles

NSX provides the system-generated IPSec tunnel profile and an IKE profile that are assigned by default when you configure either an IPSec VPN or L2 VPN service. A system-generated DPD profile is created for an IPSec VPN configuration.

The IKE and IPSec profiles provide information about the algorithms that are used to authenticate, encrypt, and establish a shared secret between network sites. The DPD profile provides information about the number of seconds to wait in between probes to detect if an IPSec peer site is alive or not.

If you decide not to use the default profiles provided by NSX, you can configure your own profile using the information in the topics that follow in this section.

Add IKE Profiles

The Internet Key Exchange (IKE) profiles provide information about the algorithms that are used to authenticate, encrypt, and establish a shared secret between network sites when you establish an IKE tunnel.

NSX provides system-generated IKE profiles that are assigned by default when you configure an IPSec VPN or L2 VPN service. The following table lists the default profiles provided.

Table 7-6. Default IKE Profiles Used for IPSec VPN or L2 VPN Services

Default IKE Profile Name	Description
nsx-default-l2vpn-ike-profile	■ Used for an L2 VPN service configuration. ■ Configured with IKE V2, AES CBC 128 encryption algorithm, SHA2 256 algorithm, and Diffie-Hellman group14 key exchange algorithm.
nsx-default-l3vpn-ike-profile	■ Used for an IPSec VPN service configuration. ■ Configured with IKE V2, AES CBC 128 encryption algorithm, SHA2 256 algorithm, and Diffie-Hellman group 14 key exchange algorithm.

Instead of the default IKE profiles used, you can also select one of the compliance suites supported starting with NSX 2.5. See About Supported Compliance Suites for more information.

If you decide not to use the default IKE profiles or compliance suites provided, you can configure your own IKE profile using the following steps.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Networking > VPN** and then click the **Profiles** tab.

3 Select the **IKE Profiles** profile type, and click **Add IKE Profile**.

4 Enter a name for the IKE profile.

5 From the **IKE Version** drop-down menu, select the IKE version to use to set up a security association (SA) in the IPSec protocol suite.

Table 7-7. IKE Versions

IKE Version	Description
IKEv1	When selected, the IPSec VPN initiates and responds to an IKEv1 protocol only.
IKEv2	This version is the default. When selected, the IPSec VPN initiates and responds to an IKEv2 protocol only.
IKE-Flex	If this version is selected and if the tunnel establishment fails with the IKEv2 protocol, the source site does not fall back and initiate a connection with the IKEv1 protocol. Instead, if the remote site initiates a connection with the IKEv1 protocol, then the connection is accepted.

6 Select the encryption, digest, and Diffie-Hellman group algorithms from the drop-down menus. You can select multiple algorithms to apply or deselect any selected algorithms you do not want to be applied.

Table 7-8. Algorithms Used

Type of Algorithm	Valid Values	Description
Encryption	■ AES 128 (default) ■ AES 256 ■ AES GCM 128 ■ AES GCM 192 ■ AES GCM 256	The encryption algorithm used during the Internet Key Exchange (IKE) negotiation. The AES 128 and AES 256 algorithms use the CBC mode of operation. The AES-GCM algorithms are supported when used with IKEv2. They are not supported when used with IKEv1.
Digest	■ SHA2 256 (default) ■ SHA1 ■ SHA2 384 ■ SHA2 512	The secure hashing algorithm used during the IKE negotiation. If AES-GCM is the only encryption algorithm selected in the **Encryption Algorithm** text box, then no hash algorithms can be specified in the **Digest Algorithm** text box, per section 8 in RFC 5282. In addition, the Psuedo-Random Function (PRF) algorithm PRF-HMAC-SHA2-256 is implicitly selected and used in the IKE security association (SA) negotiation. The PRF-HMAC-SHA2-256 algorithm must also be configured on the peer gateway in order for the phase 1 of the IKE SA negotiation to succeed. If more algorithms are specified in the **Encryption Algorithm** text box, in addition to the AES-GCM algorithm, then one or more hash algorithms can be selected in the **Digest Algorithm** text box. In addition, the PRF algorithm used in the IKE SA negotiation is implicitly determined based on the hash algorithms configured. At least one of the matching PRF algorithms must also be configured on the peer gateway in order for the phase 1 of the IKE SA negotiation to succeed. For example, if the **Encryption Algorithm** text box contains AES 128 and AES GCM 128 and SHA1 is specified in the **Digest Algorithm** text box, then the PRF-HMAC-SHA1 algorithm is used during the IKE SA negotiation. It must also be configured in the peer gateway.
Diffie-Hellman Group	■ Group 14 (default) ■ Group 2 ■ Group 5 ■ Group 15 ■ Group 16 ■ Group 19 ■ Group 20 ■ Group 21	The cryptography schemes that the peer site and the NSX Edge use to establish a shared secret over an insecure communications channel.

Note When you attempt to establish an IPSec VPN tunnel with a GUARD VPN Client (previously QuickSec VPN Client) using two encryption algorithms or two digest algorithms, the GUARD VPN Client adds additional algorithms in the proposed negotiation list. For example, if you specified AES 128 and AES 256 as the encryption algorithms and SHA2 256 and SHA2 512 as the digest algorithms to use in the IKE profile you are using to establish the IPSec VPN tunnel, the GUARD VPN Client also proposes AES 192 (using CBC mode) and SHA2 384 in the negotiation list. In this case, NSX uses the first encryption algorithm you selected when establishing the IPSec VPN tunnel.

7 Enter a security association (SA) lifetime value, in seconds, if you want it different from the default value of 86400 seconds (24 hours).

8 Provide a description and add a tag, as needed.

9 Click **Save**.

Results

A new row is added to the table of available IKE profiles. To edit or delete a non-system created profile, click the three-dot menu () and select from the list of actions available.

Add IPSec Profiles

The Internet Protocol Security (IPSec) profiles provide information about the algorithms that are used to authenticate, encrypt, and establish a shared secret between network sites when you establish an IPSec tunnel.

NSX provides system-generated IPSec profiles that are assigned by default when you configure an IPSec VPN or L2 VPN service. The following table lists the default IPSec profiles provided.

Table 7-9. Default IPSec Profiles Used for IPSec VPN or L2 VPN Services

Name of Default IPSec Profile	Description
nsx-default-l2vpn-tunnel-profile	■ Used for L2 VPN. ■ Configured with AES GCM 128 encryption algorithm and Diffie-Hellman group 14 key exchange algorithm.
nsx-default-l3vpn-tunnel-profile	■ Used for IPSec VPN. ■ Configured with AES GCM 128 encryption algorithm and Diffie-Hellman group 14 key exchange algorithm.

Instead of the default IPSec profile, you can also select one of the supported compliance suites. See About Supported Compliance Suites for more information.

If you decide not to use the default IPSec profiles or compliance suites provided, you can configure your own using the following steps.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Networking > VPN** and then click the **Profiles** tab.

3 Select the **IPSec Profiles** profile type, and click **Add IPSec Profile**.

4 Enter a name for the IPSec profile.

5 From the drop-down menus, select the encryption, digest, and Diffie-Hellman algorithms. You can select multiple algorithms to apply.

Deselect the ones you do not want used.

Table 7-10. Algorithms Used

Type of Algorithm	Valid Values	Description
Encryption	■ AES GCM 128 (default) ■ AES 128 ■ AES 256 ■ AES GCM 192 ■ AES GCM 256 ■ No Encryption Auth AES GMAC 128 ■ No Encryption Auth AES GMAC 192 ■ No Encryption Auth AES GMAC 256 ■ No Encryption	The encryption algorithm used during the Internet Protocol Security (IPSec) negotiation. The AES 128 and AES 256 algorithms use the CBC mode of operation.
Digest	■ SHA1 ■ SHA2 256 ■ SHA2 384 ■ SHA2 512	The secure hashing algorithm used during the IPSec negotiation.
Diffie-Hellman Group	■ Group 14 (default) ■ Group 2 ■ Group 5 ■ Group 15 ■ Group 16 ■ Group 19 ■ Group 20 ■ Group 21	The cryptography schemes that the peer site and NSX Edge use to establish a shared secret over an insecure communications channel.

6 Deselect **PFS Group** if you decide not to use the PFS Group protocol on your VPN service.

It is selected by default.

7 In the **SA Lifetime** text box, modify the default number of seconds before the IPSec tunnel must be re-established.

By default, an SA lifetime of 24 hours (86400 seconds) is used.

8 Select the value for **DF Bit** to use with the IPSec tunnel.

The value determines how to handle the "Don't Fragment" (DF) bit included in the data packet received. The acceptable values are described in the following table.

Table 7-11. DF Bit Values

DF Bit Value	Description
COPY	The default value. When this value is selected, NSX copies the value of the DF bit from the received packet into the packet which is forwarded. This value implies that if the data packet received has the DF bit set, after encryption, the packet also has the DF bit set.
CLEAR	When this value is selected, NSX ignores the value of the DF bit in the data packet received, and the DF bit is always 0 in the encrypted packet.

9 Provide a description and add a tag, if necessary.

10 Click **Save**.

Results

A new row is added to the table of available IPSec profiles. To edit or delete a non-system created profile, click the three-dot menu () and select from the list of actions available.

Add DPD Profiles

A DPD (Dead Peer Detection) profile provides information about the number of seconds to wait in between probes to detect if an IPSec peer site is alive or not.

NSX provides a system-generated DPD profile, named `nsx-default-l3vpn-dpd-profile`, that is assigned by default when you configure an IPSec VPN service. This default DPD profile is a periodic DPD probe mode.

If you decide not to use the default DPD profile provided, you can configure your own using the following steps.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Navigate to **Networking > VPN > Profiles**.

3 Select **DPD Profiles** from the **Select Profile Type** drop-down menu, and click **Add DPD Profile**.

4 Enter a name for the DPD profile.

5 From the **DPD Probe Mode** drop-down menu, select **Periodic** or **On Demand** mode.

For a periodic DPD probe mode, a DPD probe is sent every time the specified DPD probe interval time is reached.

For an on-demand DPD probe mode, a DPD probe is sent if no IPSec packet is received from the peer site after an idle period. The value in **DPD Probe Interval** determines the idle period used.

6 In the **DPD Probe Interval** text box, enter the number of seconds you want the NSX Edge node to wait before sending the next DPD probe.

   For a periodic DPD probe mode, the valid values are between 3 and 360 seconds. The default value is 60 seconds.

   For an on-demand probe mode, the valid values are between 1 and 10 seconds. The default value is 3 seconds.

   When the periodic DPD probe mode is set, the IKE daemon running on the NSX Edge sends a DPD probe periodically. If the peer site responds within half a second, the next DPD probe is sent after the configured DPD probe interval time has been reached. If the peer site does not respond, then the DPD probe is sent again after waiting for half a second. If the remote peer site continues not to respond, the IKE daemon resends the DPD probe again, until a response is received or the retry count has been reached. Before the peer site is declared to be dead, the IKE daemon resends the DPD probe up to a maximum of times specified in the **Retry Count** property. After the peer site is declared dead, the NSX Edge node then tears down the security association (SA) on the dead peer's link.

   When the on-demand DPD mode is set, the DPD probe is sent only if no IPSec traffic is received from the peer site after the configured DPD probe interval time has been reached.

7 In the **Retry Count** text box, enter the number of retries allowed.

   The valid values are between 1 and 100. The default retry count is 5.

8 Provide a description and add a tag, as needed.

9 To enable or disable the DPD profile, click the **Admin Status** toggle.

   By default, the value is set to **Enabled**. When the DPD profile is enabled, the DPD profile is used for all IPSec sessions in the IPSec VPN service that uses the DPD profile.

10 Click **Save**.

Results

A new row is added to the table of available DPD profiles. To edit or delete a non-system created profile, click the three-dot menu () and select from the list of actions available.

Add an Autonomous Edge as an L2 VPN Client

You can use L2 VPN to extend your Layer 2 networks to a site that is not managed by NSX. An autonomous NSX Edge can be deployed on the site, as an L2 VPN client. The autonomous NSX Edge is simple to deploy, easily programmable, and provides high-performance VPN. The autonomous NSX Edge is deployed using an OVF file on a host that is not managed by NSX. You can also enable high availability (HA) for VPN redundancy by deploying primary and secondary autonomous Edge L2 VPN clients.

Prerequisites

- Create a port group and bind it to the vSwitch on your host. Ensure that this port group accepts promiscuous mode and forged transmits from the port group's security settings.
- Create a port group for your internal L2 extension port.
- Obtain the IP addresses for the local IP and remote IP to use with the L2 VPN Client session you are adding.
- Obtain the peer code that was generated during the L2 VPN server configuration.

Procedure

1. Using vSphere Web Client, log in to the VMware vCenter that manages the non-NSX environment.
2. Select **Hosts and Clusters** and expand clusters to show the available hosts.
3. Right-click the host where you want to install the autonomous NSX Edge and select **Deploy OVF Template**.
4. Enter the URL to download and install the OVF file from the Internet or click **Browse** to locate the folder on your computer that contains the autonomous NSX Edge OVF file and click **Next**.
5. On the **Select name and folder** page, enter a name for the autonomous NSX Edge and select the folder or data center where you want to deploy. Then click **Next**.
6. On the **Select a compute resource** page, select the destination of the compute resource.
7. On the OVF Template Details page, review the template details and click **Next**.
8. On the **Configuration** page, select a deployment configuration option.
9. On the **Select storage** page, select the location to store the files for the configuration and disk files.
10. On the **Select networks** page, configure the networks that the deployed template must use. Select the port group you created for the uplink interface, the port group that you created for the L2 extension port, and enter an HA interface. Click **Next**.
11. On the **Customize Template** page, enter the following values and click **Next**.
    a. Type and retype the CLI admin password.
    b. Type and retype the CLI enable password.
    c. Type and retype the CLI root password.
    d. Enter the IPv4 address for the Management Network.
    e. Enable the option to deploy an autonomous Edge.
    f. Enter the **External Port** details for VLAN ID, exit interface, IP address, and IP prefix length such that the exit interface maps to the Network with the port group of your uplink interface.

If the exit interface is connected to a trunk port group, specify a VLAN ID. For example, **`20,eth2,192.168.5.1,24`**. You can also configure your port group with a VLAN ID and use VLAN 0 for the **External Port**.

g (Optional) To configure High Availability, enter the **HA Port** details where the exit interface maps to the appropriate HA Network.

h (Optional) When deploying an autonomous NSX Edge as a secondary node for HA, select **Deploy this autonomous-edge as a secondary node**.

Use the same OVF file as the primary node and enter the primary node's IP address, user name, password, and thumbprint.

To retrieve the thumbprint of the primary node, log in to the primary node and run the following command:

```
get certificate api thumbprint
```

Ensure that the VTEP IP addresses of the primary and secondary nodes are in the same subnet and that they connect to the same port group. When you complete the deployment and start the secondary-edge, it connects to the primary node to form an edge-cluster .

12 On the **Ready to complete** page, review the autonomous Edge settings and click **Finish**.

Note If there are errors during the deployment, a message of the day is displayed on the CLI. You can also use an API call to check for errors:

```
GET https://<nsx-mgr>/api/v1/node/status
```

The errors are categorized as soft errors and hard errors. Use API calls to resolve the soft errors as required. You can clear the message of day using an API call:

```
POST /api/v1/node/status?action=clear_bootup_error
```

13 Power on the autonomous NSX Edge appliance.

14 Log in to the autonomous NSX Edge client.

15 Select **L2VPN > Add Session** and enter the following values:

a Enter a session name.

b Enter the local IP address and the remote IP address.

c Enter the peer code from the L2VPN server. See Download the Remote Side L2 VPN Configuration File for details on obtaining the peer code.

16 Click **Save**.

17 Select **Port > Add Port** to create an L2 extension port.

18 Enter a name, a VLAN, and select an exit interface.

19 Click **Save**.

20 Select **L2VPN > Attach Port** and enter the following values:

a Select the L2 VPN session that you created.

b Select the L2 extension port that you created.

c Enter a tunnel ID.

21 Click **Attach**.

You can create additional L2 extension ports and attach them to the session if you need to extend multiple L2 networks.

22 Use the browser to log in to the autonomous NSX Edge or use API calls to view the status of the L2VPN session.

Note If the L2VPN server configuration changes, ensure that you download the peer code again and update the session with the new peer code.

Check the Realized State of an IPSec VPN Session

After you send a configuration update request for an IPSec VPN session, you can check to see if the requested state has been successfully processed in the NSX local control plane on the transport nodes.

When you create an IPSec VPN session, multiple entities are created: IKE profile, DPD profile, tunnel profile, local endpoint, IPSec VPN service, and IPSec VPN session. These entities all share the same `IPSecVPNSession` span, so you can obtain the realization state of all the entities of the IPSec VPN session by using the same `GET` API call. You can check the realization state using only the API.

Prerequisites

- Familiarize yourself with IPSec VPN. See Understanding IPSec VPN.
- Verify the IPSec VPN is configured successfully. See Add an IPSec VPN Service.
- You must have access to the NSX Manager API.

Procedure

1 Send a `POST`, `PUT`, or `DELETE` request API call.

For example:

```
PUT https://<nsx-mgr>/api/v1/vpn/ipsec/sessions/8dd1c386-9b2c-4448-85b8-51ff649fae4f
{
   "resource_type": "PolicyBasedIPSecVPNSession",
   "id": "8dd1c386-9b2c-4448-85b8-51ff649fae4f",
   "display_name": "Test RZ_UPDATED",
```

```
   "ipsec_vpn_service_id": "7adfa455-a6fc-4934-a919-f5728957364c",
   "peer_endpoint_id":  "17263ca6-dce4-4c29-bd8a-e7d12bd1a82d",
   "local_endpoint_id": "91ebfa0a-820f-41ab-bd87-f0fb1f24e7c8",
   "enabled": true,
   "policy_rules": [
       {
           "id": "1026",
           "sources": [
               {
                   "subnet": "1.1.1.0/24"
               }
           ],
           "logged": true,
           "destinations": [
               {
                   "subnet": "2.1.4..0/24"
               }
           ],
           "action": "PROTECT",
           "enabled": true,
           "_revision": 1
       }
   ]
}
```

2 Locate and copy the value of `x-nsx-requestid` from the response header returned.

For example:

```
x-nsx-requestid   e550100d-f722-40cc-9de6-cf84d3da3ccb
```

3 Request the realization state of the IPSec VPN session using the following `GET` call.

```
GET https://<nsx-mgr>/api/v1/vpn/ipsec/sessions/<ipsec-vpn-session-id>/state?request_id=<request-id>
```

The following API call uses the `id` and `x-nsx-requestid` values in the examples used in the previous steps.

```
GET https://<nsx-mgr>/api/v1/vpn/ipsec/sessions/8dd1c386-9b2c-4448-85b8-51ff649fae4f/state?
request_id=e550100d-f722-40cc-9de6-cf84d3da3ccb
```

Following is an example of a response you receive when the realization state is `in_progress`.

```
{
  "details": [
    {
      "sub_system_type": "TransportNode",
      "sub_system_id": "fe651e63-04bd-43a4-a8ec-45381a3b71b9",
      "state": "in_progress",
      "failure_message": "CCP Id:ab5958df-d98a-468e-a72b-d89dcdae5346, Message:State
realization is in progress at the node."
    },
    {
      "sub_system_type": "TransportNode",
```

```
      "sub_system_id": "ebe174ac-e4f1-4135-ba72-3dd2eb7099e3",
      "state": "in_sync"
    }
  ],
  "state": "in_progress",
  "failure_message": "The state realization is in progress at transport nodes."
}
```

Following is an example of a response you receive when the realization state is `in_sync`.

```
{
    "details": [
        {
            "sub_system_type": "TransportNode",
            "sub_system_id":  "7046e8f4-a680-11e8-9bc3-020020593f59",
            "state": "in_sync"
        }
    ],
    "state": "in_sync"
}
```

The following are examples of possible responses you receive when the realization state is `unknown`.

```
{
    "state": "unknown",
    "failure_message": "Unable to get response from any CCP node. Please retry operation
after some time."
}
```

```
{
    "details": [
       {
          "sub_system_type": "TransportNode",
          "sub_system_id": "3e643776-5def-11e8-94ae-020022e7749b",
          "state": "unknown",
          "failure_message": "CCP Id:ab5958df-d98a-468e-a72b-d89dcdae5346, Message: Unable
to get response from the node. Please retry operation after some time."
       },
       {
          "sub_system_type": "TransportNode",
          "sub_system_id": "4784ca0a-5def-11e8-93be-020022f94b73",
          "state": "in_sync"
       }
    ],
    "state": "unknown",
    "failure_message": "The state realization is unknown at transport nodes"
}
```

After you perform an entity `DELETE` operation, you might receive the status of `NOT_FOUND`, as shown in the following example.

```
{
   "http_status": "NOT_FOUND",
   "error_code": 600,
   "module_name": "common-services",
   "error_message": "The operation failed because object identifier LogicalRouter/
61746f54-7ab8-4702-93fe-6ddeb804 is missing: Object identifiers are case sensitive.."
}
```

If the IPSec VPN service associated with the session is disabled, you receive the `BAD_REQUEST` response, as shown in the following example.

```
{
    "httpStatus": "BAD_REQUEST",
    "error_code": 110199,
    "module_name": "VPN",
    "error_message": "VPN service f9cfe508-05e3-4e1d-b253-fed096bb2b63 associated with the
session 8dd1c386-9b2c-4448-85b8-51ff649fae4f is disabled. Can not get the realization
status."
}
```

Understanding TCP MSS Clamping

TCP MSS clamping enables you to reduce the maximum segment size (MSS) value used by a TCP session during a connection establishment through a VPN tunnel.

TCP MSS is the maximum amount of data in bytes that a host is willing to accept in a single TCP segment. Each end of a TCP connection sends its desired MSS value to its peer-end during the three-way handshake, where MSS is one of the TCP header options used in a TCP SYN packet. The sender host calculates the TCP MSS based on the maximum transmission unit (MTU) of its egress interface.

When a TCP traffic goes through any kind of VPN tunnel, additional headers are added to the original packet to keep it secure. For IPSec tunnel mode, additional headers used are IP, ESP, and optionally UDP (if a port translation is present in the network). Because of these additional headers, the size of the encapsulated packet goes beyond the MTU of the VPN interface. The packet can get fragmented or dropped based on the DF policy.

To avoid packet fragmentation or drop in an IPSec VPN session, you can adjust the MSS value for the IPSec session by enabling the TCP MSS clamping feature. Navigate to **Networking > VPN > IPSec Sessions**. When you are adding an IPSec session or editing an existing one, expand the **Advanced Properties** section, and enable **TCP MSS Clamping**. By default, the TCP MSS Clamping feature is disabled for an IPSec session.

When the TCP MSS Clamping feature is enabled for an IPSec session, you can configure the pre-calculated MSS value suitable for the IPSec session by setting both **TCP MSS Direction** and **TCP MSS Value**. The configured MSS value is used for MSS clamping. You can opt to use the dynamic MSS calculation by setting the **TCP MSS Direction** and leaving **TCP MSS Value** blank. The MSS value is auto-calculated based on the VPN interface MTU, VPN overhead, and the path MTU (PMTU) when it is already determined. The effective MSS is recalculated during each TCP handshake to handle the MTU or PMTU changes dynamically. See Add a Policy-Based IPSec Session or Add a Route-Based IPSec Session for more information.

Similarly, for L2 VPN, TCP MSS Clamping configuration is given only in the L2 VPN server session. You can navigate to **Networking > VPN > L2 VPN Sessions**. Select **Add L2 VPN Session > L2 VPN Server** and expand the **Advanced Properties** section. TCP MSS Clamping is enabled by default for both the directions with auto-calculation mode, but you can configure a desired TCP MSS value that is suitable for the topology or disable it. See Add an L2 VPN Server Session for more information.

Troubleshooting VPN Problems

This section provides information to help you resolve problems you might encounter while using the VPN features in NSX.

Monitor and Troubleshoot VPN Sessions

After you configure an IPSec or L2 VPN session, you can monitor the VPN tunnel status and troubleshoot any reported tunnel issues using the NSX Manager user interface.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Navigate to the **Networking > VPN > IPSec Sessions** or **Networking > VPN > L2 VPN Sessions** tab.
3. Expand the row for the VPN session that you want to monitor or troubleshoot.
4. To view the status of the VPN tunnel status, click the info icon.

   The Status dialog box appears and displays the available statuses.
5. To view the VPN tunnel traffic statistics, click **View Statistics** in the Status column.

   The Statistics dialog box displays the traffic statistics for the VPN tunnel.
6. To view the error statistics, click the **View More** link in the Statistics dialog box.
7. To close the **Statistics** dialog box, click **Close**.

Alarms When an IPsec VPN Session or Tunnel Is Down

When an IPsec VPN session or tunnel is down, an alarm is raised and the reason for the `Down` alarm is displayed on the Alarms dashboard or the VPN page on the NSX Manager user interface.

Solution

Use the following tables to locate the `Reason` message that you see on the NSX Manager user interface and review the possible cause for the `Down` alarm. To resolve the alarm, perform the necessary actions listed for the specific `Reason` message and possible cause for the `Down` alarm.

Table 7-12. Causes and Solutions for an IPsec VPN Session Is Down Alarm

Reason for the IPsec VPN Session Down Alarm	Possible Cause	Necessary Actions to Resolve the Alarm Message
`Authentication failed`	The IKE SA establishment between the VPN gateways failed due to a failure in authentication. Authentication of the IKE SA depends on the pre-shared key, Local ID, and Remote ID values.	■ Verify the `Local ID` and `Remote ID` values. The Local ID value must be set as the Remote ID value in the peer VPN gateway. ■ Verify the `Pre-shared Key` value. It must match exactly in both the VPN gateways.
`No proposal chosen`	The IKE transform configuration in both the local and peer configuration file are inconsistent.	Ensure that the following properties are configured the same for both gateways. ■ `DH groups` ■ `Digest and encryption algorithms`
`Peer sent delete`	The peer gateway initiated a delete case. A `DELETE` payload is received for IKE SA.	To determine why the peer gateway sent a `DELETE` payload, examine the logs in the NSX Edge and on the peer gateway side.
`Peer not responding`	The IKE SA negotiation timed out.	■ Verify that the remote gateway is up. ■ Verify the connectivity to the remote gateway.
`Invalid syntax`	■ IKE proposals or transforms are not formed correctly. ■ There are malformed IKE payloads.	To debug the invalid syntax, analyze the logs.
`Invalid spi`	An invalid SPI value was received in the IKE payload.	To debug the invalid SPI value, analyze the logs.
`Configuration failed`	The session configuration realization failed in NSX Edge due to some constraints or certain criteria. The reason is listed in the session dump under `Session_Config_Fail_Reason`.	Resolve the error using the reason displayed in the session dump under `Session_Config_Fail_Reason`.
`Negotiation not started`	The IKE SA negotiation has not started.	■ Verify that the `Connection Initiation Mode` property in the session configuration is set to `Responder` ■ Verify that the peer configuration has at least one of the gateways set to `Initiator`. .

Table 7-12. Causes and Solutions for an IPsec VPN Session Is Down Alarm (continued)

Reason for the IPsec VPN Session Down Alarm	Possible Cause	Necessary Actions to Resolve the Alarm Message
`IPsec service not active`	Status of the VPN service used for the session is not active.	Verify if the Admin Status in the IPsec VPN service configuration is disabled.
`Session disabled`	Admin has disabled the session.	Enable the session to resolve this error.
`SR state is not Active`	SR is in a standby state.	Verify that the VPN session on the NSX Edge node where HA peer SR is in the Active state.

Table 7-13. Causes and Solutions for an IPsec VPN Tunnel Is Down Alarm

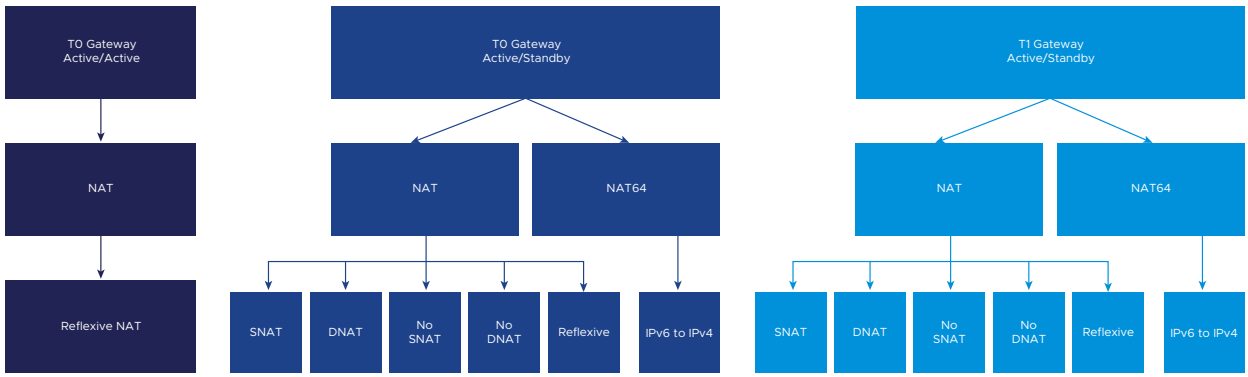
Reason for the IPsec VPN Tunnel Down Alarm	Possible Cause	Necessary Actions to Resolve the Alarm Message
`Peer sent delete`	The peer gateway sent a `DELETE` payload for the IPSEC SA.	To understand why the peer gateway sent a `DELETE` payload, you must check the logs in both the NSX Edge and in the peer gateway side.
`No proposal chosen`	The ESP transform configuration is not consistent in the configurations for both the local and peer gateways.	Ensure that the following properties are configured the same for both gateways. ■ `DH groups` ■ `Digest and encryption algorithms` ■ The `PFS` is activated or not.
`TS unacceptable`	The IPsec SA setup has failed due to a mismatch in the policy rule definition between the gateways for the tunnel configuration.	Check the local and remote network configuration on both gateways.
`IKE SA down`	The IKE SA session is down.	Check the session down reason listed in the logs and resolve the errors.
`Invalid syntax`	■ The proposals or transforms are not formed correctly. ■ There are malformed payloads.	To debug the invalid syntax, analyze the logs.
`Invalid spi`	An invalid SPI value was received in the ESP payload.	To debug the invalid SPI value, analyze the logs.
`No IKE peers`	All IKE peers are dead. There are no peer gateways left with whom to try to establish a connection.	■ Check if the remote gateway is up. ■ Check the connectivity to the configured peer gateways.
`IPsec negotiation not started`	The IPsec SA negotiation has not started.	The IKE SA is not up yet. Check the session down reason listed in the logs and resolve the errors.

Network Address Translation (NAT)

8

Network address translation (NAT) maps one IP address space to another. You can configure NAT on tier-0 and tier-1 gateways.

The following diagram shows how NAT can be configured.



Three types of NAT are supported, in addition to NAT64.

Note Disabling gateway firewall causes the NAT rule to drop traffic. If the gateway firewall needs to be disabled, include an Allow rule:

Source	Destination	ACTION
ANY	ANY	ALLOW

- Source NAT (SNAT) - translates a source IP address of outbound packets so that packets appear as originating from a different network. Supported on tier-0/tier-1 gateways running in active-standby mode. For one-to-one SNAT, the SNAT translated IP address is not programmed on the loopback port, and there is no forwarding entry with an SNAT translated IP as the prefix. For n-to-one SNAT, the SNAT translated IP address is programmed on the loopback port, and users will see a forwarding entry with an SNAT-translated IP address prefix. NSX SNAT is designed to be applied to traffic that egresses the NSX environment.
- Destination NAT (DNAT) - translates the destination IP address of inbound packets so that packets are delivered to a target address into another network. Supported on tier-0/tier-1 gateways running in active-standby mode. NSX DNAT is designed to be applied to traffic that ingresses the NSX environment.

- Reflexive NAT - (sometimes called stateless NAT) translates addresses passing through a routing device. Inbound packets undergo destination address rewriting, and outbound packets undergo source address rewriting. It is not keeping a session as it is stateless. Supported on tier-0 gateways running in active-active or active-standby mode, and on tier-1 gateways. Stateful NAT is not supported in active-active mode.

You can also disable SNAT or DNAT for an IP address or a range of addresses (No-SNAT/No-DNAT). If an address has multiple NAT rules, the rule with the highest priority is applied.

Note If there is a service interface configured in a NAT rule, `translated_port` will be realized on NSX Manager as `destination_port`. This means that the service will be the translated port while the translated port is used to match the traffic as destination port. If there is no service configured, the port will be ignored.

If you are creating a NAT rule from a Global Manager in an NSX Federation environment, you can select site-specific IP addresses for NAT. Note the following:

- Do not click **Set** under **Apply To** if you want the default option of applying the NAT rule to all locations.
- Under **Apply To**, click **Set** and select the locations whose entities you want to apply the rule to and then select **Apply NAT rule to all entities**.
- Under **Apply To**, click **Set**, select a location and then select **Interfaces** from the **Categories** drop-down menu. You can select specific interfaces to which you want to apply the NAT rule.
- DNAT is not supported on a tier-1 gateway where policy-based IPSec VPN is configured.
- SNAT configured on a tier-0 gateway's external interface processes traffic from a tier-1 gateway, and from another external interface on the tier-0 gateway.
- NAT is configured on the uplinks of the tier-0/tier-1 gateways and processes traffic going through this interface. This means that tier-0 gateway NAT rules will not apply between two tier-1 gateways connected to the tier-0.

NAT Support Matrices

Configuration Fields

Type	source-addr	dest-addr	translated-addr	translated-port	match-service
SNAT	optional	optional	must	no	optional
DNAT	optional	must	must	optional	optional
NO_SNAT	must	optional	no	no	optional
NO_DNAT	optional	must	no	no	optional
REFLEXIVE	must	no	must	no	no
NAT64	optional	must	must	optional	optional

Configuration Use Cases

Types	1:1	n:n	n:m	n:1	1:m
SNAT	Yes	Yes	Yes	Yes	No
DNAT	Yes	Yes	* configurable, but not supported	Yes	* configurable, but not supported
NO_SNAT	-	-	-	-	-
NO_DNAT	-	-	-	-	-
REFLEXIVE	Yes	Yes	No	No	No
NAT64	Yes	Yes	No	Yes	No

This chapter includes the following topics:

- Configure NAT/DNAT/No SNAT/No DNAT/Reflexive NAT
- Configure NAT64
- NAT and Gateway Firewall

Configure NAT/DNAT/No SNAT/No DNAT/Reflexive NAT

You can configure different types of NAT for IPv4 on a tier-0 or tier-1 gateway.

Note If there is a service configured in this NAT rule, the translated_port will be realized on NSX Manager as the destination_port. This means the service will be the translated port while the translated port is used to match the traffic as destination port. If there is no service configured, the port will be ignored.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Networking** > **NAT**.

3 Select a gateway from the **Gateway** drop-down menu.

4 Next to **View**, select **NAT**.

5 Click **Add NAT Rule**.

6 Enter a **Name**.

7 Select an action.

Gateway	Available Actions
Tier-1 gateway	Available actions are **SNAT**, **DNAT**, **Reflexive**, **NO SNAT**, and **NO DNAT**.
Tier-0 gateway in active-standby mode	Available actions are **SNAT**, **DNAT**, **NO SNAT**, and **NO DNAT**.
Tier-0 gateway in active-active mode	The available action is **Reflexive**.

8 Enter a **Source**. If this text box is left blank, the NAT rule applies to all sources outside of the local subnet.

   Specify an IP address, or an IP address range in CIDR format. For **SNAT**, **NO_SNAT** and **Reflexive** rules, this is a required field and represents the source network of the packets leaving the network.

9 (Required) Enter a **Destination**.

   Specify an IP address, or an IP address range in CIDR format. For **DNAT** and **NO_DNAT** rules, this is a required field and represents the source network of the packets leaving the network. This field is not applicable for **Reflexive**.

10 Enter a value for **Translated IP**.

   Specify an IPv4 address, or an IP address range in CIDR format. If translated IP is less than the match IP for SNAT it will work as PAT.

11 Toggle **Enable** to enable the rule.

12 In the **Service** column, click **Set** to select services.

   If there is a service interface configured in a NAT rule, `translated_port` will be realized on NSX Manager as `destination_port`. This means that the service will be the translated port while the translated port is used to match the traffic as destination port. If there is no service configured, the port will be ignored.

13 Enter a value for **Translated Port**.

   If there is a service interface configured in a NAT rule, `translated_port` will be realized on NSX Manager as `destination_port`. This means that the service will be the translated port while the translated port is used to match the traffic as destination port. If there is no service configured, the port will be ignored.

14 For **Apply To**, click **Set** and select objects that this rule applies to.

The available objects are **Tier-0 Gateways**, **Interfaces**, **Labels**, **Service Instance Endpoints**, and **Virtual Endpoints**.

Note If you are using NSX Federation and creating a NAT rule from a Global Manager appliance, you can select site-specific IP addresses for NAT. You can apply the NAT rule to any of the following location spans:

- Do not click **Set** if you want to use the default option of applying the NAT rule to all locations.
- Click **Set**. In the **Applied To | New Rule** dialog box, select the locations whose entities you want to apply the rule to and then click **Apply**.
- Click **Set**. In the **Applied To | New Rule** dialog box, select a location and then select **Interfaces** from the **Categories** drop-down menu. You can select specific interfaces to which you want to apply the NAT rule.
- Click **Set**. In the **Applied To | New Rule** dialog box, select a location and then select **VTI** from the **Categories** drop-down menu. You can select specific VTIs to which you want to apply the NAT rule.

See Features and Configurations Supported in NSX Federation for more details.

15 (Optional) Select a firewall setting.

The available settings are:

- **Match External Address** - The firewall will be applied to external address of a NAT rule.
  - For SNAT, the external address is the translated source address after NAT is done.
  - For DNAT, the external address is the original destination address before NAT is done.
  - For REFLEXIVE, to egress traffic, the firewall is applied to the translated source address after NAT is done. For ingress traffic, the firewall is applied to the original destination address before NAT is done.
- **Match Internal Address** - Indicates the firewall will be applied to internal address of a NAT rule.
  - For SNAT, the internal address is the original source address before NAT is done.
  - For DNAT, the internal address is the translated destination address after NAT is done.
  - For REFLEXIVE, for egress traffic, the firewall is applied to the original source address before NAT is done. For ingress traffic, the firewall is applied to the translated destination address after NAT is done.
- **Bypass** - The packet bypasses firewall rules.

16 (Optional) Toggle the **Logging** button to enable logging.

17 Specify a priority value.

A lower value means a higher priority. The default is 0. A **No SNAT** or **No DNAT** rule should have a higher priority than other rules.

18 (Optional) **Apply to Policy Based VPN**: Applies only for the **DNAT** or **No DNAT** rule category. The rule is applied based on the priority value. Despite the **Bypass** or **Match** settings, the settings applied for the **Apply To** parameter of a NAT policy are still honored.

- **Bypass:** NAT Rule is not applied to the traffic decrypted from a policy-based IPSec VPN tunnel. This is the default setting.
- **Match**: If the traffic is decrypted from a policy-based IPSec VPN tunnel, the NAT policy is evaluated and matched. NAT policy is NOT evaluated on the traffic that is not decrypted from a policy-based IPSec VPN tunnel.

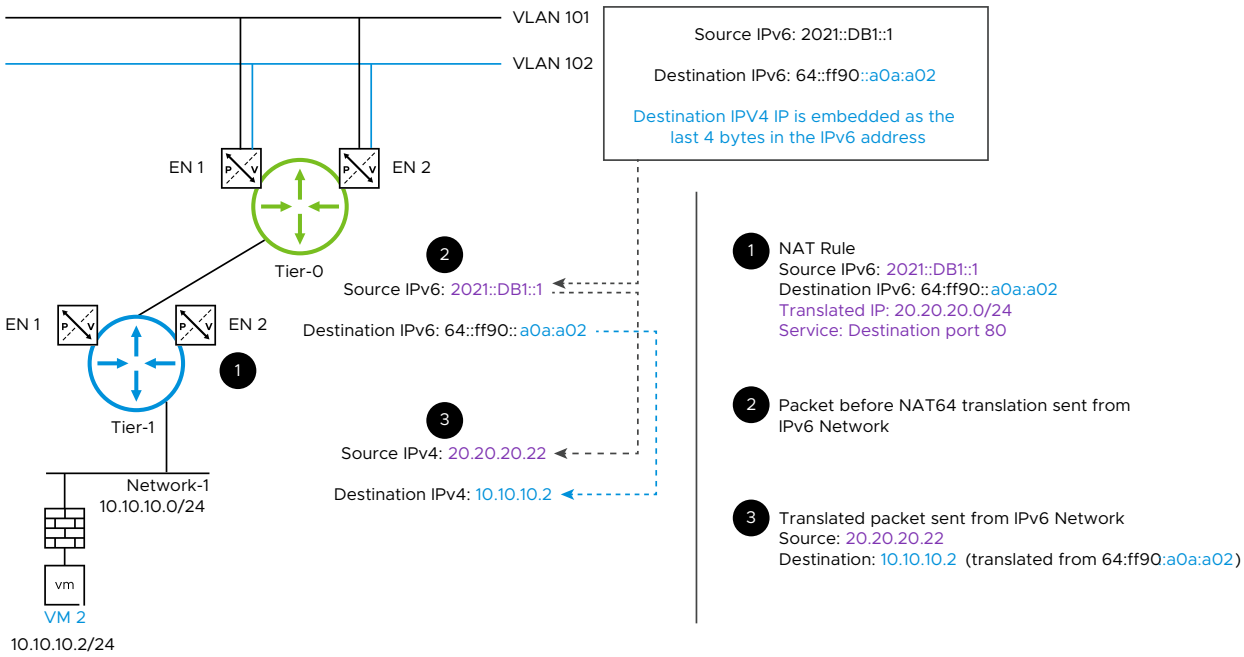
For a NAT policy to hit the decrypted traffic, the policy must be set to **Match** and the interface where encrypted traffic is sent/received must be set in the **Apply To** parameter of the NAT policy. For more information on the **Apply To** parameter, see Chapter 8 Network Address Translation (NAT).

19 Click **Save**.

Configure NAT64

NAT64 is a mechanism for translating IPv6 packets to IPv4 packets, and vice versa. NAT64 allows IPv6-only clients to contact IPv4 servers using unicast UDP or TCP. NAT64 only allows an IPv6-only client to initiate communications to an IPv4-only server. To perform IPv6-to-IPv4 translation, binding and session information is saved. NAT64 is stateful.

The following diagram shows details of NAT64 translation.



Note

- NAT64 is only supported for external IPv6 traffic coming in through the NSX Edge uplink to the IPv4 server in the overlay.
- NAT64 supports TCP and UDP. Packets of all other protocol types are discarded. NAT64 does not support ICMP, fragmentation, or IPV6 packets that have extension headers.
- When a NAT64 rule and an NSX load balancer are configured on the same Edge node, using the NAT64 rule to direct IPv6 packets to the IPv4 load balancer is not supported.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Networking** > **NAT**.

3 Select a gateway from the **Gateway** dropdown list.

4 Next to **View**, select **NAT64**.

5 Click **Add NAT 64 Rule**.

6 Enter a **Name**.

7 Enter a **Source**.

Specify an IPv6 address, or an IPv6 address range in CIDR format. For example, 2001:DB7:1::1 or 2001:DB7:1::/64.

If this text box is left blank, the NAT rule applies to all sources outside of the local subnet.

8 Enter a **Destination**.

   Specify an IPv6 address, or an IPv6 address range in CIDR format with subnet size 96. For example, 64:ff9b::0B01:0101 or 2001:DB8::/96.

9 Enter a value for **Translated IP**.

   Specify an IPv4 address, an IPv4 address range, or a comma-separated list of IPv4 addresses. For example, 10.1.1.1, 10.1.1.1-10.1.1.2, or 10.1.1.1,10.1.1.2.

10 Toggle **Enable** to enable the rule.

11 (Optional) In the **Service** column, click **Set** to select services.

12 (Optional) Enter a value for **Translated Port**.

13 (Optional) For **Apply To**, click **Set** and select objects that this rule applies to.

   The only option available is **Interfaces**.

14 (Optional) Toggle the **logging** button to enable logging.

15 (Optional) Specify a priority value.

   A lower value means a higher priority. The default is 0.

16 The **Firewall** setting is set to **Bypass** and cannot be changed.

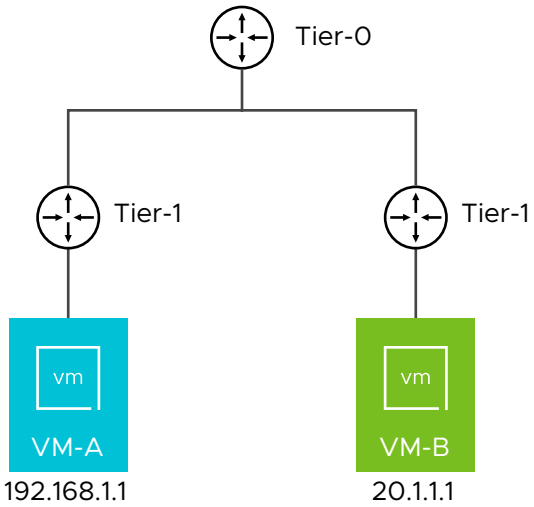
17 Click **Save**.

NAT and Gateway Firewall

A NAT firewall allows internet traffic to pass through the gateway if a device on the private network requested it. Any unsolicited requests or data packets are discarded, preventing communication with potentially dangerous devices.

If a tier-1 gateway has both SNAT and gateway firewall (GWFW) configured, and if the GWFW is not configured to be stateful, you must configure NO SNAT for the tier-1 gateway's advertised subnets. Otherwise, traffic to IP addresses in these subnets will fail.

In the example below, T1-A is the gateway and there is a SNAT rule configured that translates any traffic from its attached subnet 192.168.1.0/0 to 10.1.1.1.



Here are some traffic scenarios:

1 Any traffic stream that is initiated from VM-A/192.168.1.1 will get translated to 10.1.1.1 as the source IP, regardless if gateway firewall is stateful, stateless, or disabled. When the traffic from VM-C or VM-B returns for that flow, they will have a destination IP of 10.1.1.1; T1-A will match it up with the SNAT flow and translate it correctly so that it flows back to VM-A. The SNAT rule works as expected, and there are no issues.

2 VM-B/20.1.1.1 initiates a traffic flow to VM-A/192.168.1.1. Here, there's a difference in behavior when T1-A has a stateful firewall versus when it has no firewall or stateless firewall. The firewall rules permit the traffic between VM-B and VM-A. To have this scenario, configure a NO-NAT rule for traffic matching 192.168.1.0/24 to 20.1.1.0/24. When this NO-NAT rule exists, then there will be no difference in behavior.

3 If T1-A has stateful firewall, the T1-A firewall will create a firewall connection entry for the TCP SYN packet from VM-B/20.1.1.1 to VM-A/192.168.1.1. When VM-A replies, T1-A will match the reply packet with the stateful connection entry, and forward the traffic from VM-A/192.168.1.1 to VM-B/ 20.1.1.1 with no SNAT translation. This is because the firewall will skip the SNAT lookup when the return traffic matches up with a firewall connection entry.

4 If T1-A has firewall disabled or stateless, the T1-A firewall will forward the TCP SYN packet from VM-B/20.1.1.1 to VM-A/192.168.1.1 without creating a firewall connection entry, because it's either stateless or no firewall. When VM-A/192.168.1.1 replies back to VM-B/20.1.1.1, T1-A sees that there's no firewall connection entry, performs SNAT on it, and translates the source IP from VM-A/192.168.1.1 to 10.1.1.1. When that reply gets back to VM-B, VM-B will drop the traffic because the source IP address is 10.1.1.1 instead of VM-A/192.168.1.1.

NSX Advanced Load Balancer (Avi)

9

The VMware NSX Advanced Load Balancer is a distributed and highly scalable cloud-native Application Distribution solution.

Starting in 3.2.2, the NSX Advanced Load Balancer is no longer consumed on the NSX Manager. If you have activated the NSX Advanced Load Balancer, or have upgraded from 3.2.0 or 3.2.1 to 3.2.2 or higher, we recommend you deactivate the NSX Advanced Load Balancer by clicking **Deactivate NSX-T ALB** in the banner message on the UI.

After deactivation none of the configurations are lost, and there is no impact on the running traffic. The configurations are preserved on the NSX Advanced Load Balancer.

After deactivating the NSX Advanced Load Balancer from the NSX Manager, log into the NSX Advanced Load Balancer Controller to access all of the NSX Advanced Load Balancer configurations.

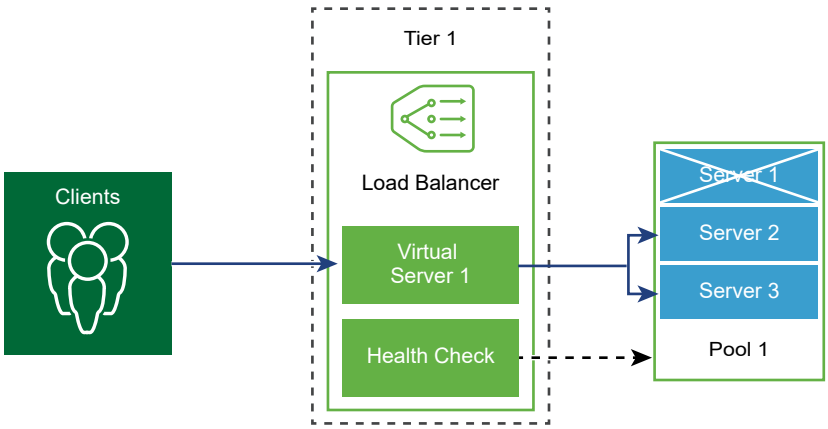
To install and configure a controller cluster see Install NSX Advanced Load Balancer Appliance Cluster, in the *NSX Installation Guide*.

For NSX Advanced Load Balancer configuration, see VMware NSX Advanced Load Balancer Documentation.

Load Balancer

10

The NSX logical load balancer offers high-availability service for applications and distributes the network traffic load among multiple servers.



The load balancer distributes incoming service requests evenly among multiple servers in such a way that the load distribution is transparent to users. Load balancing helps in achieving optimal resource utilization, maximizing throughput, minimizing response time, and avoiding overload.

You can map a virtual IP address to a set of pool servers for load balancing. The load balancer accepts TCP, UDP, HTTP, or HTTPS requests on the virtual IP address and decides which pool server to use.

Depending on your environment needs, you can scale the load balancer performance by increasing the existing virtual servers and pool members to handle heavy network traffic load.

Note Logical load balancer is supported only on the tier-1 gateway. One load balancer can be attached only to a tier-1 gateway.

This chapter includes the following topics:

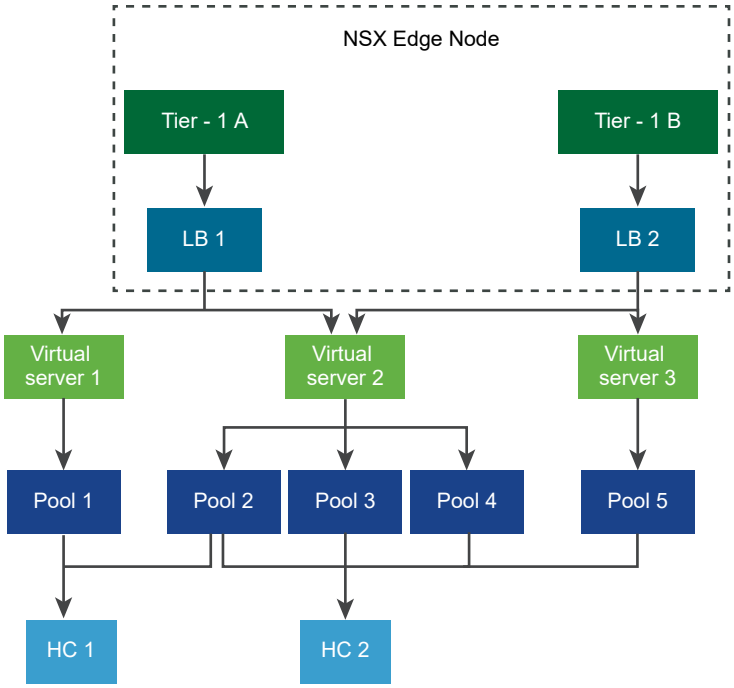
- Key Load Balancer Concepts
- Setting Up Load Balancer Components
- Groups Created for Server Pools and Virtual Servers

Key Load Balancer Concepts

Load balancer includes virtual servers, server pools, and health checks monitors.

A load balancer is connected to a Tier-1 logical router. The load balancer hosts single or multiple virtual servers. A virtual server is an abstract of an application service, represented by a unique combination of IP, port, and protocol. The virtual server is associated to single to multiple server pools. A server pool consists of a group of servers. The server pools include individual server pool members.

To test whether each server is correctly running the application, you can add health check monitors that check the health status of a server.



Scaling Load Balancer Resources

When you configure a load balancer, you can specify a size (small, medium, large, or extra large). The size determines the number of virtual servers, server pools, and pool members the load balancer can support.

A load balancer runs on a tier-1 gateway, which must be in active-standby mode. The gateway runs on NSX Edge nodes. The form factor of the NSX Edge node (bare metal, small, medium, large, or extra large) determines the number of load balancers that the NSX Edge node can support. Note that in Manager mode, you create logical routers, which have similar functionality to gateways. See Chapter 1 NSX Manager.

For more information about what the different load balance sizes and NSX Edge form factors can support, see https://configmax.vmware.com.

Note that using a small NSX Edge node to run a small load balancer is not recommended in a production environment.

You can call an API to get the load balancer usage information of an NSX Edge node. If you use Policy mode to configure load balancing, run the following command:

```
GET /policy/api/v1/infra/lb-node-usage?node_path=<node-path>
```

If you use Manager mode to configure load balancing, run the following command:

```
GET /api/v1/loadbalancer/usage-per-node/<node-id>
```

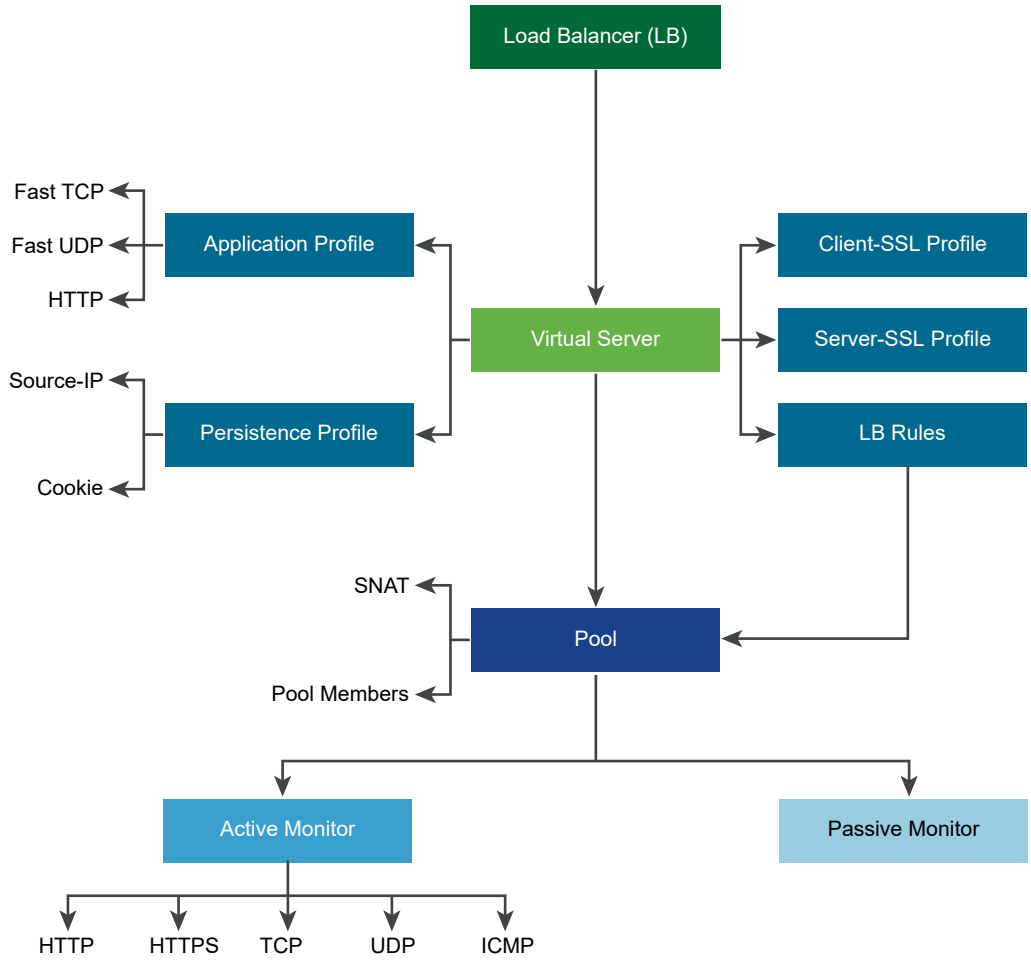
The usage information includes the number of load balancer objects (such as load balancer services, virtual servers, server pools, and pool members) that are configured on the node. For more information, see the *NSX API Guide*.

Supported Load Balancer Features

NSX load balancer supports the following features.

- Layer 4 - TCP and UDP
- Layer 7 - HTTP and HTTPS with load balancer rules support
- Server pools - static and dynamic with NSGroup
- Persistence - Source-IP and Cookie persistence mode
- Health check monitors - Active monitor which includes HTTP, HTTPS, TCP, UDP, and ICMP, and passive monitor
- SNAT - Transparent, Automap, and IP List
- HTTP upgrade - For applications using HTTP upgrade such as WebSocket, the client or server requests for HTTP Upgrade, which is supported. By default, NSX supports and accepts HTTPS upgrade client request using the HTTP application profile.

Note: SSL -Terminate-mode and proxy-mode is not supported in NSX limited export release.

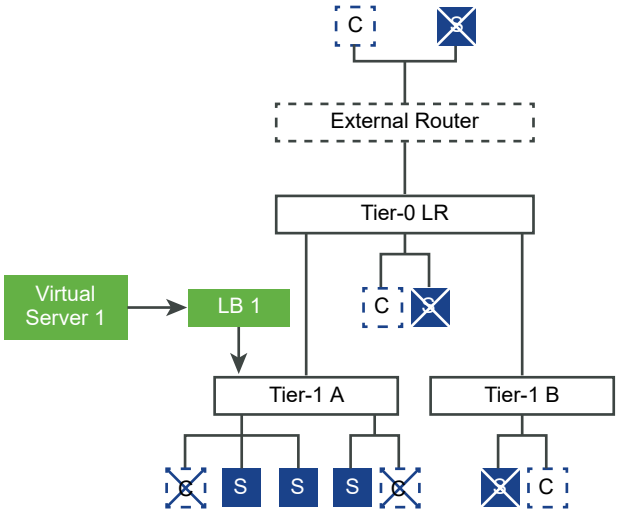


Load Balancer Topologies

Load balancers are typically deployed in either inline or one-arm mode. One-arm mode requires virtual server Source NAT (SNAT) configuration, and inline mode does not.

Inline Topology

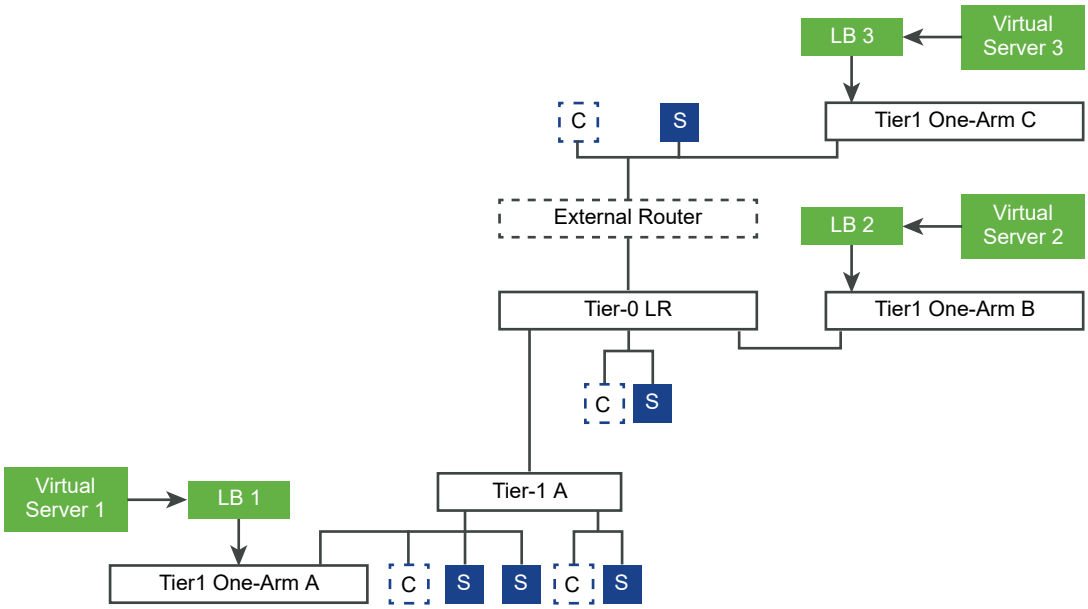
In the inline mode, the load balancer is in the traffic path between the client and the server. Clients and servers should not be connected to overlay segments on the same tier-1 logical router if SNAT on the load balancer is not desired. If clients and servers are connected to overlay segments on the same tier-1 logical router, SNAT is required.



One-Arm Topology

In one-arm mode, the load balancer is not in the traffic path between the client and the server. In this mode, the client and the server can be anywhere. The load balancer performs Source NAT (SNAT) to force return traffic from the server destined to the client to go through the load balancer. This topology requires virtual server SNAT to be enabled.

When the load balancer receives the client traffic to the virtual IP address, the load balancer selects a server pool member and forwards the client traffic to it. In the one-arm mode, the load balancer replaces the client IP address with the load balancer IP address so that the server response is always sent to the load balancer. The load balancer forwards the response to the client.



Special Use Case When no overlay is used and everything is VLAN, overlay still must be configured on the edge nodes hosting the tier-1 one-arm load balancer. This is because edge nodes must have at least one tunnel end point UP for high availability between edge nodes. When the tunnels are UP they will agree on which edge node will run the active or standby role of each tier-0 and tier-1 gateway.

Tier-1 Service Chaining

If a tier-1 gateway or logical router hosts different services, such as NAT, firewall, and load balancer, the services are applied in the following order:

- Ingress

  DNAT - Firewall - Load Balancer

  Note: If DNAT is configured with Firewall Bypass, firewall is skipped but not load balancer.

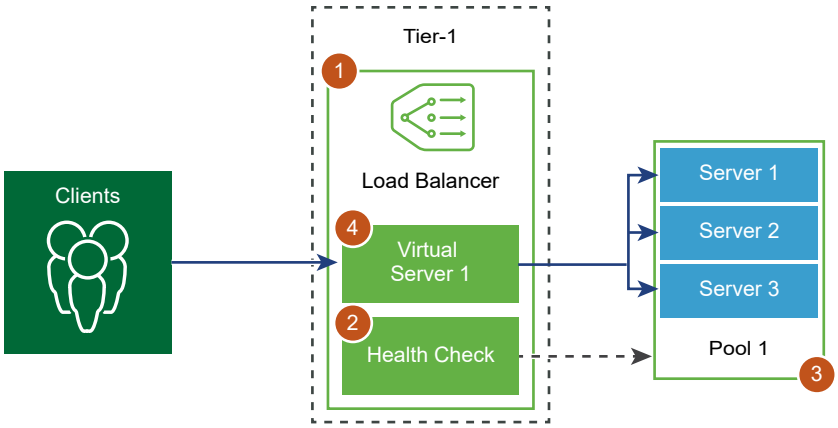
- Egress

  Load Balancer - Firewall - SNAT

Setting Up Load Balancer Components

To use logical load balancers, you must start by configuring a load balancer and attaching it to a tier-1 gateway.

Next, you set up health check monitoring for your servers. You must then configure server pools for your load balancer. Finally, you must create a layer 4 or layer 7 virtual server for your load balancer and attach the newly created virtual server to the load balancer.



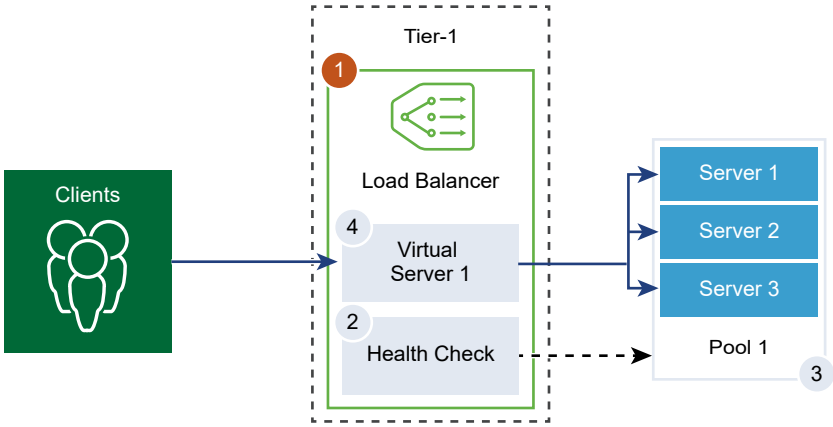
Add Load Balancers

Load balancer is created and attached to the tier-1 gateway.

You can configure the level of error messages you want the load balancer to add to the error log.

Note

- Avoid setting the log level to DEBUG on load balancers with a significant traffic due to the number of messages printed to the log that affect performance.
- Load balancer over IPSec VPN is not supported for route-based VPN terminated on Tier-1 gateways.



Prerequisites

Verify that a tier-1 gateway is configured. See Chapter 3 Tier-1 Gateway.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Select **Networking > Load Balancing > Add Load Balancer**.
3. Enter a name and a description for the load balancer.
4. Select the load balancer size based on your needs of virtual servers and pool members and available resources.
5. Select the already configured tier-1 gateway to attach to this load balancer from the drop-down menu.

   The tier-1 gateway must be in the Active-Standby mode.
6. Define the severity level of the error log from the drop-down menu.

   Load balancer collects information about encountered issues of different severity levels to the error log.
7. (Optional) Enter tags to make searching easier.

   You can specify a tag to set a scope of the tag.

8 Click **Save**.

The load balancer creation and attaching the load balancer to the tier-1 gateway takes about three minutes and the configuration status to appear green and Up.

If the status is Down, click the information icon and resolve the error before you proceed.

9 (Optional) Delete the load balancer.

a Detach the load balancer from the virtual server and tier-1 gateway.

b Select the load balancer.

c Click the vertical ellipses button.

d Select **Delete**.

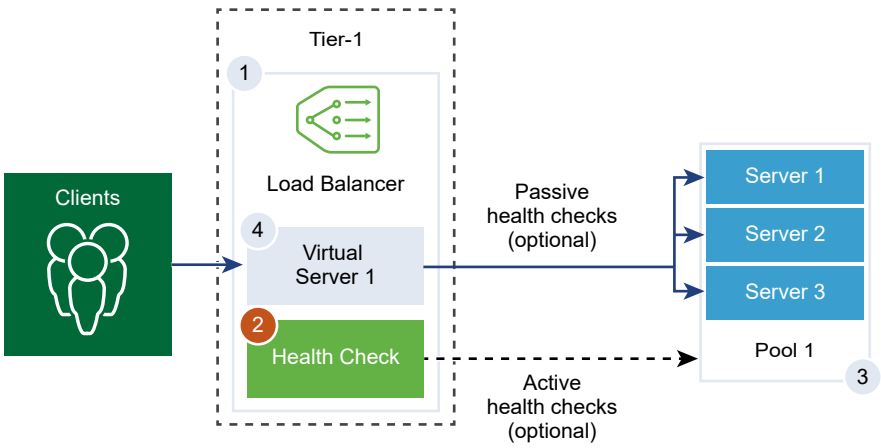
Add an Active Monitor

The active health monitor is used to test whether a server is available. The active health monitor uses several types of tests such as sending a basic ping to servers or advanced HTTP requests to monitor an application health.

Servers that fail to respond within a certain time period or respond with errors are excluded from future connection handling until a subsequent periodic health check finds these servers to be healthy.

Active health checks are performed on server pool members after the pool member is attached to a virtual server and that virtual server is attached to a tier-1 gateway. The tier-1 uplink IP address is used for the health check.

Note More than one active health monitor can be configured per server pool.



Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Networking > Load Balancing > Monitors > Active > Add Active Monitor**.

3 Select a protocol for the server from the drop-down menu.

You can also use predefined protocols; HTTP, HTTPS, ICMP, TCP, and UDP for NSX Manager.

4 Select the **HTTP** protocol.

5 Configure the values to monitor a service pool.

You can also accept the default active health monitor values.

Option	Description
Name and Description	Enter a name and description for the active health monitor.
Monitoring Port	Set the value of the monitoring port.
Monitoring Interval	Set the time in seconds that the monitor sends another connection request to the server.
Timeout Period	Set the time the load balancer will wait for the Pool Member monitor response before considering failed.
Fail Count	Set a value when the consecutive failures reach this value, the server is considered temporarily unavailable.
Rise Count	Set the value of consecutive successful monitors to reach before changing the Pool Member Status from Down to Up.
Tags	Enter tags to make searching easier. You can specify a tag to set a scope of the tag.

For example, if the monitoring interval is set as 5 seconds and the timeout as 15 seconds, the load balancer send requests to the server every 5 seconds. In each probe if the expected response is received from the server within 15 seconds, the health check result is OK. If not, then the result is CRITICAL. If the recent three health check results are all UP, the server is considered as UP.

6 To configure the HTTP Request, click **Configure**.

7 Enter the HTTP request and response configuration details.

Option	Description
HTTP Method	Select the method to detect the server status from the drop-down menu, GET, OPTIONS, POST, HEAD, and PUT.
HTTP Request URL	Enter the request URI for the method. ASCII control characters (backspace, vertical tab, horizontal tab, line feed, etc), unsafe characters such as a `space`, `\`, `<`, `>`, `{`, `}`, and any character outside the ASCII character set are not allowed in the request URL and should be encoded. For example, replace a space with a plus (+) sign, or with %20.
HTTP Request Version	Select the supported request version from the drop-down menu. You can also accept the default version, HTTP_VERSION_1.
HTTP Request Header	Click **Add** and enter the HTTP request header name and corresponding value.

Option	Description
HTTP Request Body	Enter the request body. Valid for the POST and PUT methods.
HTTP Response Code	Enter the string that the monitor expects to match in the status line of HTTP response body. The response code is a comma-separated list. For example, 200,301,302,401.
HTTP Response Body	If the HTTP response body string and the HTTP health check response body match, then the server is considered as healthy.

8 Click **Save**.

9 Select the **HTTPS** protocol from the drop-down list.

10 Complete step 5.

11 Click **Configure**.

12 Enter the HTTP request and response and SSL configuration details.

Option	Description
Name and Description	Enter a name and description for the active health monitor.
HTTP Method	Select the method to detect the server status from the drop-down menu, GET, OPTIONS, POST, HEAD, and PUT.
HTTP Request URL	Enter the request URI for the method. ASCII control characters (backspace, vertical tab, horizontal tab, line feed, etc), unsafe characters such as a `space`, `\`, `<`, `>`, `{`, `}`, and any character outside the ASCII character set are not allowed in the request URL and should be encoded. For example, replace a space with a plus (+) sign, or with %20.
HTTP Request Version	Select the supported request version from the drop-down menu. You can also accept the default version, HTTP_VERSION_1.
HTTP Request Header	Click **Add** and enter the HTTP request header name and corresponding value.
HTTP Request Body	Enter the request body. Valid for the POST and PUT methods.
HTTP Response Code	Enter the string that the monitor expects to match in the status line of HTTP response body. The response code is a comma-separated list. For example, 200,301,302,401.
HTTP Response Body	If the HTTP response body string and the HTTP health check response body match, then the server is considered as healthy.
Server SSL	Toggle the button to enable the SSL server.
Client Certificate	(Optional) Select a certificate from the drop-down menu to be used if the server does not host multiple host names on the same IP address or if the client does not support an SNI extension.

Option	Description
Server SSL Profile	(Optional) Assign a default SSL profile from the drop-down menu that defines reusable and application-independent client-side SSL properties. Click the vertical ellipses and create a custom SSL profile.
Trusted CA Certificates	(Optional) You can require the client to have a CA certificate for authentication.
Mandatory Server Authentication	(Optional) Toggle the button to enable server authentication.
Certificate Chain Depth	(Optional) Set the authentication depth for the client certificate chain.
Certificate Revocation List	(Optional) Set a Certificate Revocation List (CRL) in the client-side SSL profile to reject compromised client certificates.

13 Select the **ICMP** protocol.

14 Complete step 5 and assign the data size in byte of the ICMP health check packet.

15 Select the **TCP** protocol.

16 Complete step 5 and you can leave the TCP data parameters empty.

If both the data sent and expected are not listed, then a three-way handshake TCP connection is established to validate the server health. No data is sent.

Expected data if listed has to be a string. Regular expressions are not supported.

17 Select the **UDP** protocol.

18 Complete step 5 and configure the UDP data.

Required Option	Description
UDP Data Sent	Enter the string to be sent to a server after a connection is established.
UDP Data Expected	Enter the string expected to receive from the server. Only when the received string matches this definition, is the server is considered as UP.

What to do next

Associate the active health monitor with a server pool. See Add a Server Pool.

Add a Passive Monitor

Load balancers perform passive health checks to monitor failures during client connections and mark servers causing consistent failures as DOWN.

Passive health check monitors client traffic going through the load balancer for failures. For example, if a pool member sends a TCP Reset (RST) in response to a client connection, the load balancer detects that failure. If there are multiple consecutive failures, then the load balancer considers that server pool member to be temporarily unavailable and stops sending

connection requests to that pool member for some time. After some time, the load balancer sends a connection request to verify that the pool member has recovered. If that connection is successful, then the pool member is considered healthy. Otherwise, the load balancer waits for some time and tries again.

Passive health check considers the following scenarios to be failures in the client traffic.

- For server pools associated with Layer 7 virtual servers, if the connection to the pool member fails. For example, if the pool member sends a TCP RST when the load balancer tries to connect or perform an SSL handshake between the load balancer and the pool member fails.
- For server pools associated with Layer 4 TCP virtual servers, if the pool member sends a TCP RST in response to client TCP SYN or does not respond at all.
- For server pools associated with Layer 4 UDP virtual servers, if a port is unreachable or a destination unreachable ICMP error message is received in response to a client UDP packet.

Server pools associated to Layer 7 virtual servers, the failed connection count is incremented when any TCP connection errors, for example, TCP RST failure to send data or SSL handshake failures occur.

Server pools associated to Layer 4 virtual servers, if no response is received to a TCP SYN sent to the server pool member or if a TCP RST is received in response to a TCP SYN, then the server pool member is considered as DOWN. The failed count is incremented.

For Layer 4 UDP virtual servers, if an ICMP error such as, port or destination unreachable message is received in response to the client traffic, then it is considered as DOWN.

Note One passive health monitor can be configured per server pool.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Networking > Load Balancing > Monitors > Passive > Add Passive Monitor**.

3 Enter a name and description for the passive health monitor.

4 Configure the values to monitor a service pool.

You can also accept the default active health monitor values.

Option	Description
Fail Count	Set a value when the consecutive failures reach this value, the server is considered temporarily unavailable.
Timeout Period	Set the number of times the server is tested before it is considered as DOWN.
Tags	Enter tags to make searching easier. You can specify a tag to set a scope of the tag.

For example, when the consecutive failures reach the configured value 5, that member is considered temporarily unavailable for 5 seconds. After this period, that member is tried again for a new connection to see if it is available. If that connection is successful, then the member is considered available and the failed count is set to zero. However, if that connection fails, then it is not used for another timeout interval of 5 seconds.

What to do next

Associate the passive health monitor with a server pool. See Add a Server Pool.

Add a Server Pool

A server pool consists of one or more servers that are configured and running the same application. A single pool can be associated to both Layer 4 and Layer 7 virtual servers.

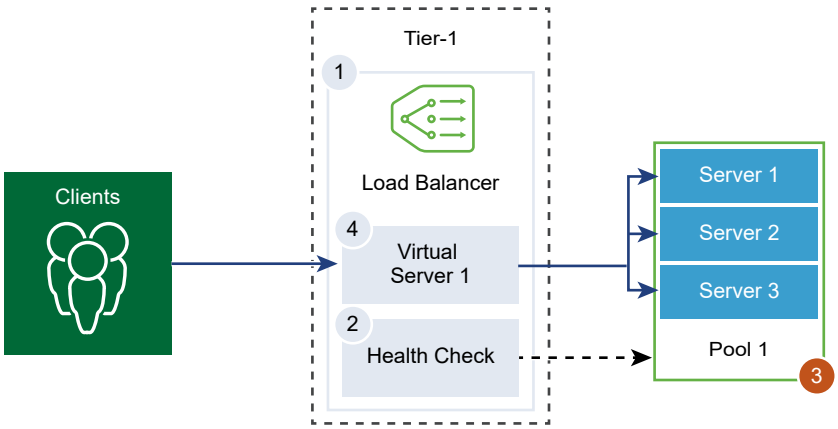
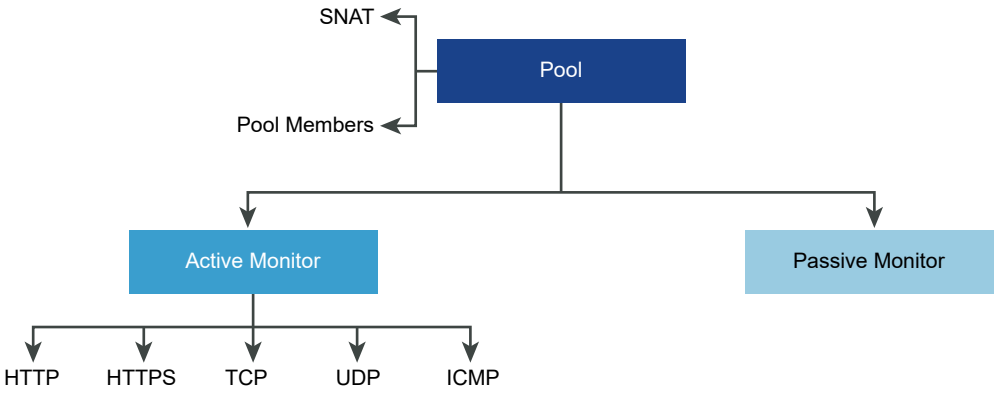


Figure 10-1. Server Pool Parameter Configuration



Prerequisites

- If you use dynamic pool members, a NSGroup must be configured. See Create an NSGroup in Manager Mode.
- Verify that a passive health monitor is configured. See Add a Passive Monitor.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Networking > Load Balancing > Server Pools > Add Server Pool**.

3 Enter a name and description for the load balancer server pool.

You can optionally describe the connections managed by the server pool.

4 Select the algorithm balancing method for the server pool.

Load balancing algorithm controls how the incoming connections are distributed among the members. The algorithm can be used on a server pool or a server directly.

All load balancing algorithms skip servers that meet any of the following conditions:

- Admin state is set to DISABLED.
- Admin state is set to GRACEFUL_DISABLED and no matching persistence entry.
- Active or passive health check state is DOWN.
- Connection limit for the maximum server pool concurrent connections is reached.

Option	Description
ROUND_ROBIN	Incoming client requests are cycled through a list of available servers capable of handling the request. Ignores the server pool member weights even if they are configured.
WEIGHTED_ROUND_ROBIN	Each server is assigned a weight value that signifies how that server performs relative to other servers in the pool. The value determines how many client requests are sent to a server compared to other servers in the pool. This load balancing algorithm focuses on fairly distributing the load among the available server resources.
LEAST_CONNECTION	Distributes client requests to multiple servers based on the number of connections already on the server. New connections are sent to the server with the fewest connections. Ignores the server pool member weights even if they are configured.
WEIGHTED_LEAST_CONNECTION	Each server is assigned a weight value that signifies how that server performs relative to other servers in the pool. The value determines how many client requests are sent to a server compared to other servers in the pool. This load balancing algorithm focuses on using the weight value to distribute the load among the available server resources. By default, the weight value is 1 if the value is not configured and slow start is enabled.
IP-HASH	Selects a server based on a hash of the source IP address and the total weight of all the running servers.

5 Click **Select Members** and elect the server pool members.

A server pool consists of single or multiple pool members.

Option	Description
Enter individual members	Enter a pool member name, IPv4 or IPv6 address, and a port. IP addresses can be either IPv4 or IPv6. Mixed addressing is not supported. Note that the pool members IP version must match the VIP IP version. For example, VIP-IPv4 with Pool-IPv4, and IPv6 with Pool-IPv6. Each server pool member can be configured with a weight for use in the load balancing algorithm. The weight indicates how much more or less load a given pool member can handle relative to other members in the same pool. You can set the server pool admin state. By default, the option is enable when a server pool member is added. If the option is disabled, active connections are processed, and the server pool member is not selected for new connections. New connections are assigned to other members of the pool. If gracefully disabled, it allows you to remove servers for maintenance. The existing connections to a member in the server pool in this state continue to be processed. Toggle the button to designate a pool member as a backup member to work with the health monitor to provide an Active-Standby state. Traffic failover occurs for backup members if active members fail a health check. Backup members are skipped during the server selection. When the server pool is inactive, the incoming connections are sent to only the backup members that are configured with a sorry page indicating an application is unavailable. **Max Concurrent Connection** value assigns a connection maximum so that the server pool members are not overloaded and skipped during server selection. If a value is not specified, then the connection is unlimited.
Select a group	Select a pre-configured group of server pool members. Enter a group name and an optional description. Set the compute member from existing list or create one. You can specify membership criteria, select members of the group, add IP addresses, and MAC addresses as group members, and add Active Directory groups. IP addresses can be either IPv4 or IPv6. Mixed addressing is not supported. The identity members intersect with the compute member to define membership of the group. Select a tag from the drop-down menu. You can optionally define the maximum group IP address list.

6 Click **Set Monitors** and select one or more active health check monitors for the server. Click **Apply**.

The load balancer periodically sends an ICMP ping to the servers to verify health independent of data traffic. You can configure more than one active health check monitor per server pool.

7 Select the Source NAT (SNAT) translation mode.

Depending on the topology, SNAT might be required so that the load balancer receives the traffic from the server destined to the client. SNAT can be enabled per server pool.

SNAT Translation Mode	Description
Automap Mode	Load Balancer uses the interface IP address and ephemeral port to continue the communication with a client initially connected to one of the server's established listening ports. SNAT is required. Enable port overloading to allow the same SNAT IP and port to be used for multiple connections if the tuple (source IP, source port, destination IP, destination port, and IP protocol) is unique after the SNAT process is performed. You can also set the port overload factor to allow the maximum number of times a port can be used simultaneously for multiple connections.
Disabled	Disable the SNAT translation mode.
IP Pool	Specify a single IPv4 or IPv6 address range, for example, 1.1.1.1-1.1.1.10 to be used for SNAT while connecting to any of the servers in the pool. IP addresses can be either IPv4 or IPv6. Mixed addressing is not supported. By default, the port range 4096 through 65535 is used for all configured SNAT IP addresses. The port range 1000 through 4095 is reserved for purposes such as health checks, and connections initiated from Linux applications. If multiple IP addresses are present, then they are selected in a Round Robin manner. If a virtual server IP port is in the SNAT default port range 4096 through 65535, make sure that the virtual server IP is **not** in the SNAT IP pool. Enable port overloading to allow the same SNAT IP and port to be used for multiple connections if the tuple (source IP, source port, destination IP, destination port, and IP protocol) is unique after the SNAT process is performed. You can also set the port overload factor to allow the maximum number of times a port can be used simultaneously for multiple connections.

8 Click **Additional Properties**, and toggle the button to enable TCP Multiplexing.

With TCP multiplexing, you can use the same TCP connection between a load balancer and the server for sending multiple client requests from different client TCP connections.

9 Set the **Max Multiplexing Connections** per server that are kept alive to send future client requests.

10 Enter the **Min Active Members** the server pool must always maintain.

11 Select a passive health monitor for the server pool from the drop-down menu.

12 Select a tag from the drop-down menu.

Setting Up Virtual Server Components

You can set up the Layer 4 and Layer 7 virtual servers and configure several virtual server components such as, application profiles, persistent profiles, and load balancer rules.

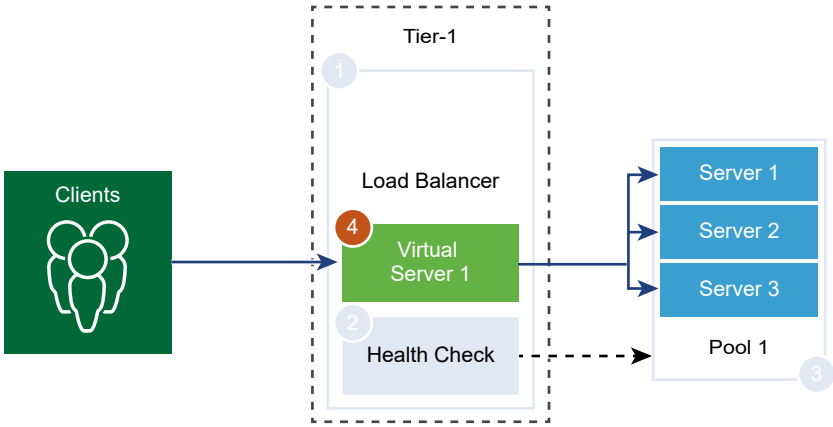
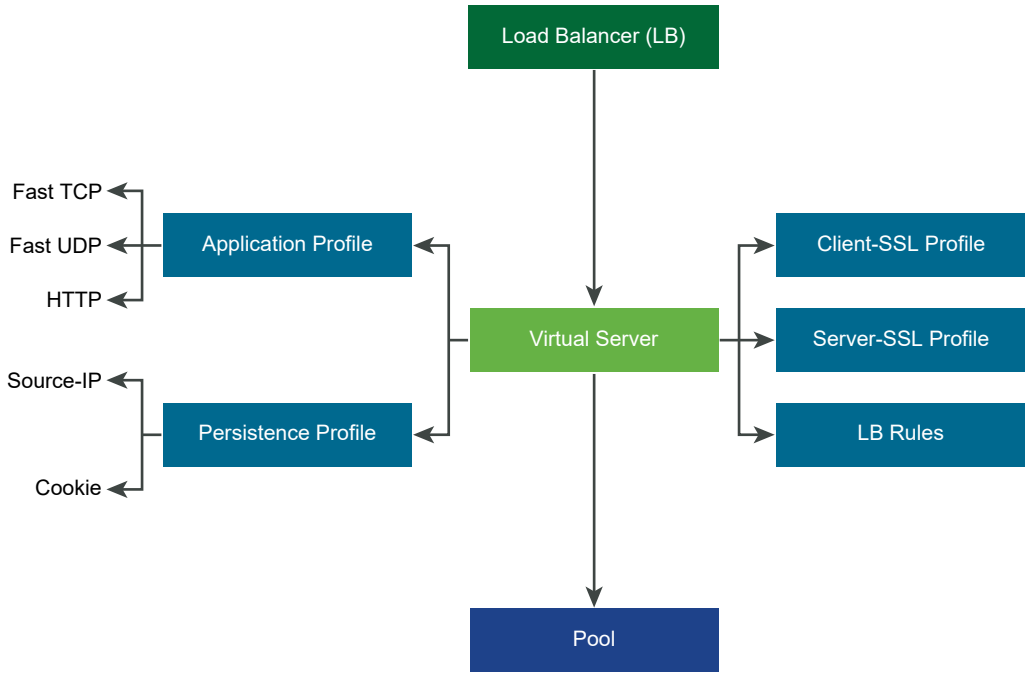


Figure 10-2. Virtual Server Components



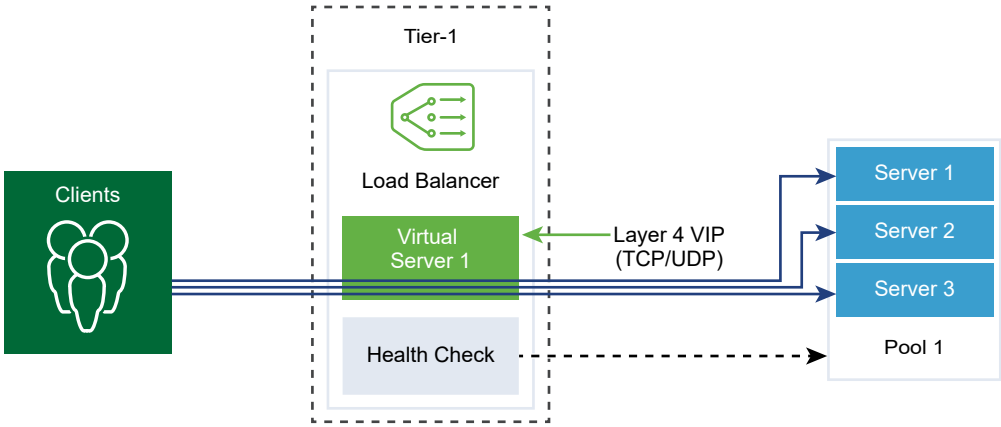
Add an Application Profile

Application profiles are associated with virtual servers to enhance load balancing network traffic and simplify traffic-management tasks.

Application profiles define the behavior of a particular type of network traffic. The associated virtual server processes network traffic according to the values specified in the application profile. Fast TCP, Fast UDP, and HTTP application profiles are the supported types of profiles.

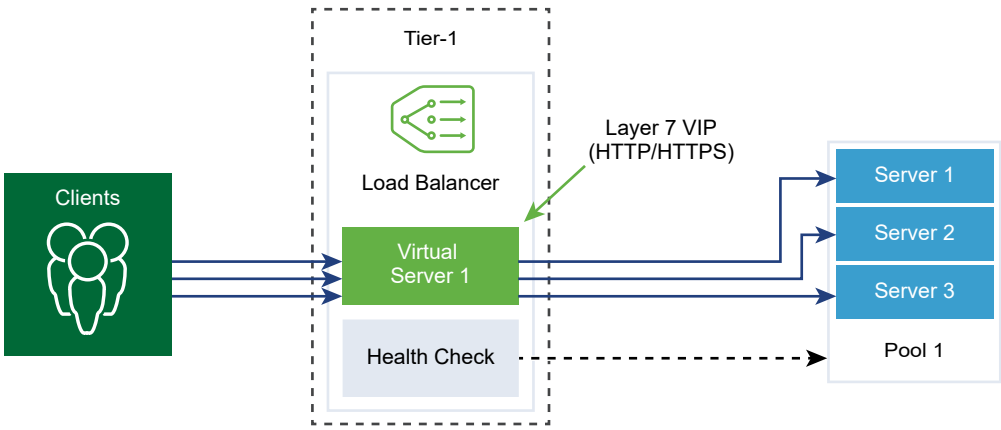
TCP application profile is used by default when no application profile is associated to a virtual server. TCP and UDP application profiles are used when an application is running on a TCP or UDP protocol and does not require any application level load balancing such as, HTTP URL load balancing. These profiles are also used when you only want Layer 4 load balancing, which has a faster performance and supports connection mirroring.

Figure 10-3. Layer 4 TCP and UDP Application Profile



HTTP application profile is used for both HTTP and HTTPS applications when the load balancer must take actions based on Layer 7 such as, load balancing all images requests to a specific server pool member or stopping HTTPS to offload SSL from pool members. Unlike the TCP application profile, the HTTP application profile terminates the client TCP connection at the load balancer and waits for the clients HTTP or HTTPS request before selecting the server pool member.

Figure 10-4. Layer 7 HTTPS Application Profile



Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Networking > Load Balancing > Profiles > Application > Add Application Profiles**.

3 Select a **Fast TCP** application profile and enter the profile details.

You can also accept the default FAST TCP profile settings.

Option	Description
Name and Description	Enter a name and a description for the Fast TCP application profile.
Idle Timeout	Enter the time in seconds on how long the server can remain idle after a TCP connection is established. Set the idle time to the actual application idle time and add a few more seconds so that the load balancer does not close its connections before the application does.
HA Flow Mirroring	Toggle the button to make all the flows to the associated virtual server mirrored to the HA standby node.
Connection Close Timeout	Enter the time in seconds that the TCP connection both FINs or RST must be kept for an application before closing the connection. A short closing timeout might be required to support fast connection rates.
Tags	Enter tags to make searching easier. You can specify a tag to set a scope of the tag.

4 Select a **Fast UDP** application profile and enter the profile details.

You can also accept the default UDP profile settings.

Option	Description
Name and Description	Enter a name and a description for the Fast UDP application profile.
Idle Timeout	Enter the time in seconds on how long the server can remain idle after a UDP connection is established. UDP is a connectionless protocol. For load balancing purposes, all the UDP packets with the same flow signature such as, source and destination IP address or ports and IP protocol received within the idle timeout period are considered to belong to the same connection and sent to the same server. If no packets are received during the idle timeout period, the connection which is an association between the flow signature and the selected server is closed.
HA Flow Mirroring	Toggle the button to make all the flows to the associated virtual server mirrored to the HA standby node.
Tags	Enter tags to make searching easier. You can specify a tag to set a scope of the tag.

5 Select a **HTTP** application profile and enter the profile details.

You can also accept the default HTTP profile settings.

To detect an inactive client or server communication, the load balancer uses the HTTP application profile response timeout feature set to 60 seconds. If the server does not send traffic during the 60 seconds interval, NSX ends the connection on the client and server side. Default application profiles cannot be edited. To edit HTTP application profile settings, create a custom profile.

HTTP application profile is used for both HTTP and HTTPS applications.

Option	Description
Name and Description	Enter a name and a description for the HTTP application profile.
Idle Timeout	Enter the time in seconds on how long client idle connections remain before the load balancer closes them (FIN).
Request Header Size	Specify the maximum buffer size in bytes used to store HTTP request headers.
Response Header Size	Specify the maximum buffer size in bytes used to store HTTP response headers. The default is 4096, and the maximum is 65536.
Redirection	■ None - If a website is temporarily down, user receives a page not found error message. ■ HTTP Redirect - If a website is temporarily down or has moved, incoming requests for that virtual server can be temporarily redirected to a URL specified here. Only a static redirection is supported. For example, if HTTP Redirect is set to http://sitedown.abc.com/sorry.html, then irrespective of the actual request, for example, http://original_app.site.com/home.html or http://original_app.site.com/somepage.html, incoming requests are redirected to the specified URL when the original website is down. ■ HTTP to HTTPS Redirect - Certain secure applications might want to force communication over SSL, but instead of rejecting non-SSL connections, they can redirect the client request to use SSL. With HTTP to HTTPS Redirect, you can preserve both the host and URI paths and redirect the client request to use SSL. For HTTP to HTTPS redirect, the HTTPS virtual server must have port 443 and the same virtual server IP address must be configured on the same load balancer. For example, a client request for http://app.com/path/page.html is redirected to https://app.com/path/page.html. If either the host name or the URI must be modified while redirecting, for example, redirect to https://secure.app.com/path/page.html, then load balancing rules must be used.
Tags	Enter tags to make searching easier. You can specify a tag to set a scope of the tag.

Option	Description
X-Forwarded-For (XFF)	■ **Insert** - If the XFF HTTP header is not present in the incoming request, the load balancer inserts a new XFF header with the client IP address. If the XFF HTTP header is present in the incoming request, the load balancer appends the XFF header with the client IP address. ■ **Replace** - If the XFF HTTP header is present in the incoming request, the load balancer replaces the header. Web servers log each request they handle with the requesting client IP address. These logs are used for debugging and analytic purposes. If the deployment topology requires SNAT on the load balancer, then server uses the client SNAT IP address which defeats the purpose of logging. As a workaround, the load balancer can be configured to insert XFF HTTP header with the original client IP address. Servers can be configured to log the IP address in the XFF header instead of the source IP address of the connection.
Request Body Size	Enter value for the maximum size of the buffer used to store the HTTP request body. If the size is not specified, then the request body size is unlimited.
Response Timeout (sec)	Enter the time in seconds on how long the load balancer waits for Server HTTP Response before it stops and closes the connection to the pool member and retries the request to another server.
Server Keep-Alive	Toggle the button for the load balancer to turn off TCP multiplexing and enable HTTP keep-alive. If the client uses HTTP/1.0, the load balancer upgrades to HTTP/1.1 protocol and the HTTP keep-alive is set. All HTTP requests received on the same client-side TCP connection are sent to the same server over a single TCP connection to ensure that reauthorization is not required. When HTTP keep-alive is enabled and forwarding rules are configured in the load balancer, the server keep-alive setting takes precedence. As a result, HTTP requests are sent to servers already connected with keep-alive. If you always want to give priority to the forwarding rules when the load balancer rule conditions are met, disable the keep-alive setting. Note that the persistence setting takes precedence over the keep-alive setting. Processing is done in the order of Persistence > Keep-Alive > Load Balancer Rules.

Add a Persistence Profile

To ensure stability of stateful applications, load balancers implement persistence which directs all related connections to the same server. Different types of persistence are supported to address different types of application needs.

Some applications maintain the server state such as, shopping carts. Such state can be per client and identified by the client IP address or per HTTP session. Applications can access or modify this state while processing subsequent related connections from the same client or HTTP session.

The source IP persistence profile tracks sessions based on the source IP address. When a client requests a connection to a virtual server that enables the source address persistence, the load balancer checks if that client was previously connected, and if so, returns the client to the same server. If not, the load balancer selects the server pool member based on the pool load balancing algorithm. Source IP persistence profile is used by Layer 4 and Layer 7 virtual servers.

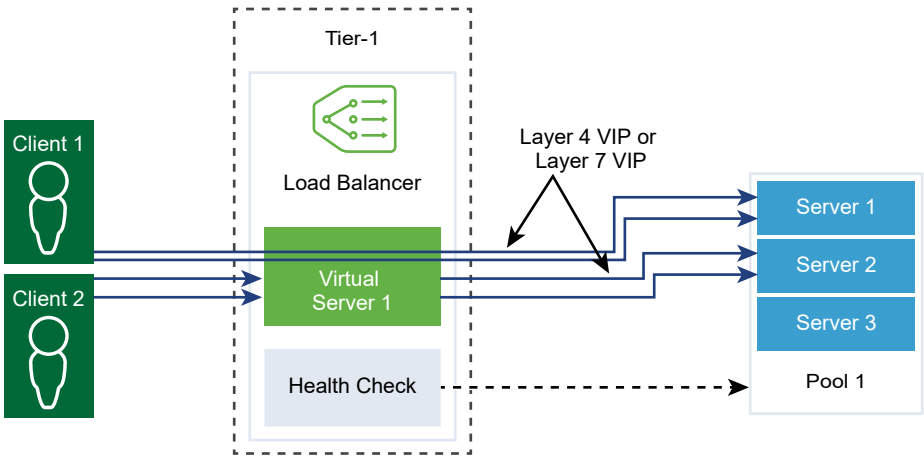
If rule persistence, cookie persistence, and server keep-alive are all configured, the load balancer follows the priority of rule persistence > cookie persistence > server keep-alive.

The Cookie persistence profile offers 3 modes:

- Cookie Insert - the load balance inserts its own cookie with the pool member information (encoded or not) in the server response to the client. The client then forwards the received cookies in subsequent requests (NSX cookie included), and the load balancer uses that information to provide the pool member persistence. The NSX cookie is trimmed from the client request when sent to the pool member.
- Cookie Prefix - the load balancer appends the pool member information (encoded or not) in the server response to the client. The client then forwards the received HTTP cookie in subsequent requests (with the NSX prepended information), and the load balancer uses that information to provide the pool member persistence. The NSX cookie prefix is trimmed from the client request when sent to the pool member.
- Cookie Rewrite - the load balancer replace server cookie value with the pool member information (encoded or not) in the server response to the client. The client then forwards the received HTTP cookie in subsequent requests (with the NSX prepended information), and the load balancer uses that information to provide the pool member persistence. The original server cookie is replaced in the client request when sent to the pool member.

Cookie persistence is available only on L7 virtual servers. Note that a blank space in a cookie name is **not** supported.

The generic persistence profile supports persistence based on the HTTP header, cookie, or URL in the HTTP request. Therefore, it supports application session persistence when the session ID is part of the URL. This profile is not associated with a virtual server directly. Specify this profile when you configure a load balancer rule for request forwarding and response rewrite.



Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Networking > Load Balancing > Profiles > Persistence > Add Persistence Profiles**.

3 Select **Source IP** to add a source IP persistence profile and enter the profile details.

You can also accept the default Source IP profile settings.

Option	Description
Name and Description	Enter a name and a description for the Source IP persistence profile.
Share Persistence	Toggle the button to share the persistence so that all virtual servers this profile is associated with can share the persistence table. If the persistence sharing is not enabled in the Source IP persistence profile associated to a virtual server, each virtual server that the profile is associated to maintains a private persistence table.
Persistence Entry Timeout	Enter the persistence expiration time in seconds. The load balancer persistence table maintains entries to record that client requests are directed to the same server. The very first connection from new client IP is load balanced to a pool member based on the load balancing algorithm. NSX will store that persistence entry on the LB persistence-table which is viewable on the Edge Node hosting the T1-LB active via the CLI command:`get load-balancer <LB-UUID> persistence-tables`. ■ When there are connections from that client to the VIP, the persistence entry is kept. ■ When there are no more connections from that client to the VIP, the persistence entry begins the timer count down specified in the "Persistence Entry Timeout" value. If no new connection from that client to the VIP is made before the timer expires, the persistence entry for that client IP is deleted. If that client comes back after the entry is deleted, it will be load balanced again to a pool member based on the load balancing algorithm.
Purge Entries When Full	A large timeout value can lead to the persistence table quickly filling up when the traffic is heavy. When this option is enabled, the oldest entry is deleted to accept the newest entry. When this option is disabled, if the source IP persistence table is full, new client connections are rejected.
HA Persistence Mirroring	Toggle the button to synchronize persistence entries to the HA peer. When HA persistence mirroring is enabled, the client IP persistence remains in the case of load balancer failover.
Tags	Enter tags to make searching easier. You can specify a tag to set a scope of the tag.

4 Select a **Cookie** persistence profile, and enter the profile details. Cookie persistence is available only on L7 virtual servers. Note that a blank space in a cookie name is **not** supported.

Option	Description
Name and Description	Enter a name and a description for the Cookie persistence profile.
Share Persistence	Toggle the button to share persistence across multiple virtual servers that are associated to the same pool members. The Cookie persistence profile inserts a cookie with the format, *<name>.<profile-id>.<pool-id>*. If the persistence shared is not enabled in the Cookie persistence profile associated with a virtual server, the private Cookie persistence for each virtual server is used and is qualified by the pool member. The load balancer inserts a cookie with the format, *<name>.<virtual_server_id>.<pool_id>*.
Cookie Mode	Select a mode from the drop-down menu. ■ INSERT - Adds a unique cookie to identify the session. ■ PREFIX - Appends to the existing HTTP cookie information. ■ REWRITE - Rewrites the existing HTTP cookie information.
Cookie Name	Enter the cookie name. A blank space in a cookie name is **not** supported.
Cookie Domain	Enter the domain name. HTTP cookie domain can be configured only in the **INSERT** mode.
Cookie Fallback	Toggle the button so that the client request is rejected if cookie points to a server that is in a DISABLED or is in a DOWN state. Selects a new server to handle a client request if the cookie points to a server that is in a DISABLED or is in a DOWN state.
Cookie Path	Enter the cookie URL path. HTTP cookie path can be set only in the **INSERT** mode.
Cookie Garbling	Toggle the button to disable encryption. When garbling is disabled, the cookie server IP address and port information is in a plain text. Encrypt the cookie server IP address and port information.
Cookie Type	Select a cookie type from the drop-down menu. **Session Cookie** - Not stored. Will be lost when the browser is closed. **Persistence Cookie** - Stored by the browser. Not lost when the browser is closed.
HttpOnly Flag	When enabled, this option prevents a script running in the browser from accessing cookies. HttpOnly Flag is only available in the **INSERT** mode.
Secure Flag	When enabled, this option causes web browsers to send cookies over https only. Secure Flag is only available in the **INSERT** mode.
Max Idle Time	Enter the time in seconds that the cookie type can be idle before a cookie expires.

Option	Description
Max Cookie Age	For the session cookie type, enter the time in seconds a cookie is available.
Tags	Enter tags to make searching easier. You can specify a tag to set a scope of the tag.

5 Select **Generic** to add a generic persistence profile and enter the profile details.

Option	Description
Name and Description	Enter a name and a description for the Source IP persistence profile.
Share Persistence	Toggle the button to share the profile among virtual servers.
Persistence Entry Timeout	Enter the persistence expiration time in seconds. The load balancer persistence table maintains entries to record that client requests are directed to the same server. The very first connection from new client IP is load balanced to a pool member based on the load balancing algorithm. NSX will store that persistence entry on the LB persistence-table which is viewable on the Edge Node hosting the T1-LB active via the CLI command:`get load-balancer <LB-UUID> persistence-tables`. ■ When there are connections from that client to the VIP, the persistence entry is kept. ■ When there are no more connections from that client to the VIP, the persistence entry begins the timer count down specified in the "Persistence Entry Timeout" value. If no new connection from that client to the VIP is made before the timer expires, the persistence entry for that client IP is deleted. If that client comes back after the entry is deleted, it will be load balanced again to a pool member based on the load balancing algorithm.
HA Persistence Mirroring	Toggle the button to synchronize persistence entries to the HA peer.
Tags	Enter tags to make searching easier. You can specify a tag to set a scope of the tag.

Add an SSL Profile

SSL profiles configure application-independent SSL properties such as, cipher lists and reuse these lists across multiple applications. SSL properties are different when the load balancer is acting as a client and as a server, as a result separate SSL profiles for client-side and server-side are supported.

Note SSL profile is not supported in the NSX limited export release.

Client-side SSL profile refers to the load balancer acting as an SSL server and stopping the client SSL connection. Server-side SSL profile refers to the load balancer acting as a client and establishing a connection to the server.

You can specify a cipher list on both the client-side and server-side SSL profiles.

SSL session caching allows the SSL client and server to reuse previously negotiated security parameters avoiding the expensive public key operation during the SSL handshake. SSL session caching is disabled by default on both the client-side and server-side.

SSL session tickets are an alternate mechanism that allows the SSL client and server to reuse previously negotiated session parameters. In SSL session tickets, the client and server negotiate whether they support SSL session tickets during the handshake exchange. If supported by both, server can send an SSL ticket, which includes encrypted SSL session parameters to the client. The client can use that ticket in subsequent connections to reuse the session. SSL session tickets are enabled on the client-side and disabled on the server-side.

Figure 10-5. SSL Offloading

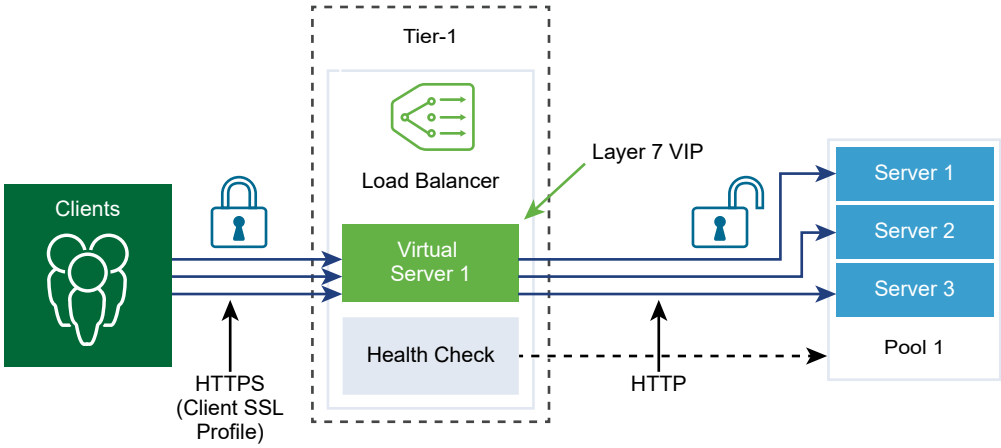
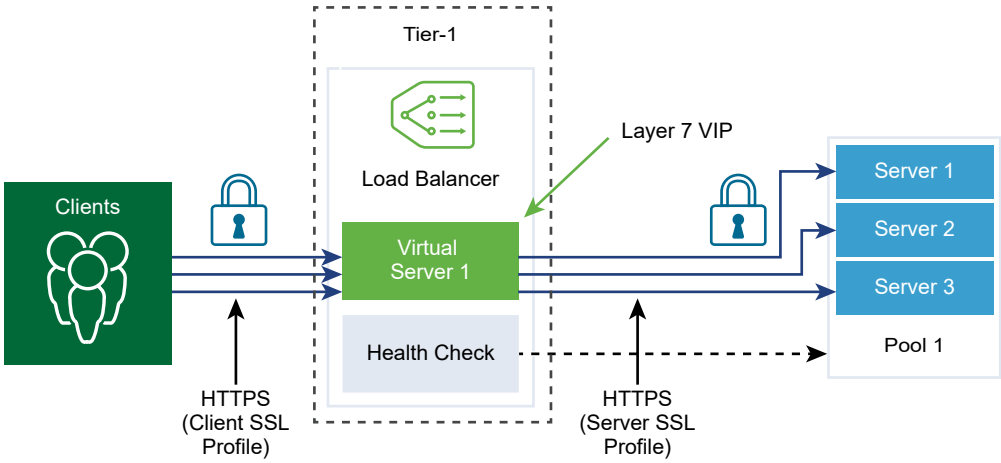


Figure 10-6. End-to-End SSL



Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Networking > Load Balancing > Profiles > SSL Profile**.

3 Select a **Client SSL Profile** and enter the profile details.

Option	Description
Name and Description	Enter a name and a description for the Client SSL profile.
SSL Suite	Select the SSL Cipher group from the drop-down menu and available SSL Ciphers and SSL protocols to be included in the Client SSL profile are populated. Balanced SSL Cipher group is the default.
Session Caching	Toggle the button to allow the SSL client and server to reuse previously negotiated security parameters avoiding the expensive public key operation during an SSL handshake.
Tags	Enter tags to make searching easier. You can specify a tag to set a scope of the tag.
Supported SSL Ciphers	Depending on the SSL suite, you assigned the supported SSL Ciphers are populated here. Click **View More** to view the entire list. If you selected **Custom**, you must select the SSL Ciphers from the drop-down menu.
Supported SSL Protocols	Depending on the SSL suite, you assigned the supported SSL protocols are populated here. Click **View More** to view the entire list. If you selected **Custom**, you must select the SSL Ciphers from the drop-down menu.
Session Cache Entry Timeout	Enter the cache timeout in seconds to specify how long the SSL session parameters must be kept and can be reused.
Prefer Server Cipher	Toggle the button so that the server can select the first supported cipher from the list it can support. During an SSL handshake, the client sends an ordered list of supported ciphers to the server.

4 Select a **Server SSL Profile** and enter the profile details.

Option	Description
Name and Description	Enter a name and a description for the Server SSL profile.
SSL Suite	Select the SSL Cipher group from the drop-down menu and available SSL Ciphers and SSL protocols to be included in the Server SSL profile are populated. Balanced SSL Cipher group is the default.
Session Caching	Toggle the button to allow the SSL client and server to reuse previously negotiated security parameters avoiding the expensive public key operation during an SSL handshake.
Tags	Enter tags to make searching easier. You can specify a tag to set a scope of the tag.
Supported SSL Ciphers	Depending on the SSL suite, you assigned the supported SSL Ciphers are populated here. Click **View More** to view the entire list. If you selected **Custom**, you must select the SSL Ciphers from the drop-down menu.

Option	Description
Supported SSL Protocols	Depending on the SSL suite, you assigned the supported SSL protocols are populated here. Click **View More** to view the entire list. If you selected **Custom**, you must select the SSL Ciphers from the drop-down menu.
Session Cache Entry Timeout	Enter the cache timeout in seconds to specify how long the SSL session parameters must be kept and can be reused.
Prefer Server Cipher	Toggle the button so that the server can select the first supported cipher from the list it can support. During an SSL handshake, the client sends an ordered list of supported ciphers to the server.

Add Layer 4 Virtual Servers

Virtual servers receive all the client connections and distribute them among the servers. A virtual server has an IP address, a port, and a protocol. For Layer 4 virtual servers, lists of ports ranges can be specified instead of a single TCP or UDP port to support complex protocols with dynamic ports.

A Layer 4 virtual server must be associated to a primary server pool, also called a default pool.

If a virtual server status is disabled, any new connection attempts to the virtual server are rejected by sending either a TCP RST for the TCP connection or ICMP error message for UDP. New connections are rejected even if there are matching persistence entries for them. Active connections continue to be processed. If a virtual server is deleted or disassociated from a load balancer, then active connections to that virtual server fail.

Prerequisites

- Verify that application profiles are available. See Add an Application Profile.
- Verify that persistent profiles are available. See Add a Persistence Profile.
- Verify that SSL profiles for the client and server are available. See Add an SSL Profile.
- Verify that server pools are available. See Add a Server Pool.
- Verify that load balancer is available. See Add Load Balancers.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Select **Networking > Load Balancing > Virtual Servers > Add Virtual Server**.
3. Select a **L4 TCP** or a **L4 UDP** protocol and enter the protocol details.

   Layer 4 virtual servers support either the Fast TCP or Fast UDP protocol, but not both.

For Fast TCP or Fast UDP protocol support on the same IP address and port, for example DNS, a virtual server must be created for each protocol.

L4 TCP Option	L4 TCP Description
Name and Description	Enter a name and a description for the Layer 4 virtual server.
IP Address	Enter the virtual server IP address. Both IPv4 and IPv6 addresses are supported. Note that the pool members IP version must match the VIP IP version. For example, VIP-IPv4 with Pool-IPv4, and IPv6 with Pool-IPv6.
Ports	Enter the virtual server port number.
Load Balancer	Select an existing load balancer to attach to this Layer 4 virtual server from the drop-down menu.
Server Pool	Select an existing server pool from the drop-down menu. The server pool consists of one or more servers, also called pool members that are similarly configured and running the same application. You can click the vertical ellipses to create a server pool.
Application Profile	Based on the protocol type, the existing application profile is automatically populated. Click the vertical ellipses to create an application profile.
Persistence	Select an existing persistence profile from the drop-down menu. Persistence profile can be enabled on a virtual server to allow Source IP related client connections to be sent to the same server.
Access List Control	When you enable Access List Control (ALC), all traffic flowing through the load balancer is compared with the ACL statement, which either drops or allows the traffic. ACL is disabled by default. To enable, click **Configure**, and select **Enabled**. Select an Action: ■ Allow - Allows connections matching the selected group. All other connections are dropped. ■ Drop - Allows connections not matching the selected group. A dropped connection generates a log entry is access log is enabled. Select a **Group**. The IP addresses included in this group are either dropped or allowed by the ACL.
Max Concurrent Connection	Set the maximum concurrent connection allowed to a virtual server so that the virtual server does not deplete resources of other applications hosted on the same load balancer.
Max New Connection Rate	Set the maximum new connection to a server pool member so that a virtual server does not deplete resources.
Sorry Server Pool	Select an existing sorry server pool from the drop-down menu. The sorry server pool serves the request when a load balancer cannot select a backend server to the serve the request from the default pool. You can click the vertical ellipses to create a server pool.

L4 TCP Option	L4 TCP Description
Default Pool Member Port	Enter a default pool member port if the pool member port for a virtual server is not defined. For example, if a virtual server is defined with a port range of 2000–2999 and the default pool member port range is set as 8000-8999, then an incoming client connection to the virtual server port 2500 is sent to a pool member with a destination port set to 8500.
Admin State	Toggle the button to disable the admin state of the Layer 4 virtual server.
Access Log	Toggle the button to enable logging for the Layer 4 virtual server.
Tags	Enter tags to make searching easier. You can specify a tag to set a scope of the tag.

L4 UDP Option	L4 UDP Description
Name and Description	Enter a name and a description for the Layer 4 virtual server.
IP Address	Enter the virtual server IP address. Both IPv4 and IPv6 addresses are supported. Note that the pool members IP version must match the VIP IP version. For example, VIP-IPv4 with Pool-IPv4, and IPv6 with Pool-IPv6.
Ports	Enter the virtual server port number.
Load Balancer	Select an existing load balancer to attach to this Layer 4 virtual server from the drop-down menu.
Server Pool	Select an existing server pool from the drop-down menu. The server pool consists of one or more servers, also called pool members that are similarly configured and running the same application. You can click the vertical ellipses to create a server pool.
Application Profile	Based on the protocol type, the existing application profile is automatically populated. You can click the vertical ellipses to create an application profile.
Persistence	Select an existing persistence profile from the drop-down menu. Persistence profile can be enabled on a virtual server to allow Source IP related client connections to be sent to the same server.
Max Concurrent Connection	Set the maximum concurrent connection allowed to a virtual server so that the virtual server does not deplete resources of other applications hosted on the same load balancer.
Access List Control	When you enable Access List Control (ALC) all traffic flowing through the load balancer will be compared with the ACL statement, which will either drop it or allow it. ACL is disabled by default. To enable, click **Configure**, and check **Enabled**. Select an Action: ■ Allow - Allows connections matching the selected group. All other connections are dropped ■ Drop - Allows connections not matching the selected group. A dropped connection generates a log entry is access log is enabled. Select a **Group**. The IP addresses included in this group are either dropped or allowed by the ACL.

L4 UDP Option	L4 UDP Description
Max New Connection Rate	Set the maximum new connection to a server pool member so that a virtual server does not deplete resources.
Sorry Server Pool	Select an existing sorry server pool from the drop-down menu. The sorry server pool serves the request when a load balancer cannot select a backend server to the serve the request from the default pool. You can click the vertical ellipses to create a server pool.
Default Pool Member Port	Enter a default pool member port if the pool member port for a virtual server is not defined. For example, if a virtual server is defined with port range 2000–2999 and the default pool member port range is set as 8000-8999, then an incoming client connection to the virtual server port 2500 is sent to a pool member with a destination port set to 8500.
Admin State	Toggle the button to disable the admin state of the Layer 4 virtual server.
Access Log	Toggle the button to enable logging for the Layer 4 virtual server.
Log Significant Event Only	This field can only be configured if access logs are enabled. Connections that cannot be sent to a pool member are treated as a significant event such as "max connection limit," or "Access Control drop."
Tags	Enter tags to make searching easier. You can specify a tag to set a scope of the tag.

Add Layer 7 HTTP Virtual Servers

Virtual servers receive all the client connections and distribute them among the servers. A virtual server has an IP address, a port, and a protocol TCP.

If a virtual server status is disabled, any new connection attempts to the virtual server are rejected by sending either a TCP RST for the TCP connection or ICMP error message for UDP. New connections are rejected even if there are matching persistence entries for them. Active connections continue to be processed. If a virtual server is deleted or disassociated from a load balancer, then active connections to that virtual server fail.

Note SSL profile is not supported in the NSX limited export release.

If a client-side SSL profile binding is configured on a virtual server but not a server-side SSL profile binding, then the virtual server operates in an SSL-terminate mode, which has an encrypted connection to the client and plain text connection to the server. If both the client-side and server-side SSL profile bindings are configured, then the virtual server operates in SSL-proxy mode, which has an encrypted connection both to the client and the server.

Associating server-side SSL profile binding without associating a client-side SSL profile binding is currently not supported. If a client-side and a server-side SSL profile binding is not associated with a virtual server and the application is SSL-based, then the virtual server operates in an SSL-unaware mode. In this case, the virtual server must be configured for Layer 4. For example, the virtual server can be associated to a fast TCP profile.

Prerequisites

- Verify that application profiles are available. See Add an Application Profile.
- Verify that persistent profiles are available. See Add a Persistence Profile.
- Verify that SSL profiles for the client and server are available. See Add an SSL Profile.
- Verify that server pools are available. See Add a Server Pool.
- Verify that CA and client certificate are available. See Chapter 23 Certificates.
- Verify that a certification revocation list (CRL) is available. See Import a Certificate Revocation List.
- Verify that load balancer is available. See Add Load Balancers.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Networking > Load Balancing > Virtual Servers > Add Virtual Server**.

3 Select **L7 HTTP** from the drop-down list and enter the protocol details.

Layer 7 virtual servers support the HTTP and HTTPS protocols.

Option	Description
Name and Description	Enter a name and a description for the Layer virtual server.
IP Address	Enter the virtual server IP address. Both IPv4 and IPv6 addresses are supported.
Ports	Enter the virtual server port number.
Load Balancer	Select an existing load balancer to attach to this Layer 4 virtual server from the drop down menu.
Server Pool	Select an existing server pool from the drop-down menu. The server pool consists of one or more servers, also called pool members that are similarly configured and running the same application. You can click the vertical ellipses to create a server pool.
Application Profile	Based on the protocol type, the existing application profile is automatically populated. You can click the vertical ellipses to create an application profile.
Persistence	Select an existing persistence profile from the drop-down menu. Persistence profile can be enabled on a virtual server to allow Source IP and Cookie related client connections to be sent to the same server.

4 Click **Configure** to set the Layer 7 virtual server SSL.

You can configure the Client SSL and Server SSL.

5 Configure the Client SSL.

Option	Description
Client SSL	Toggle the button to enable the profile. Client-side SSL profile binding allows multiple certificates, for different host names to be associated to the same virtual server.
Default Certificate	Select a default certificate from the drop-down menu. This certificate is used if the server does not host multiple host names on the same IP address or if the client does not support Server Name Indication (SNI) extension. To use a 2k/3k/4k certificate/key use NSX Manager to Creating Self-signed Certificates or to Create a Certificate Signing Request File. To use an 8k certificate/key, import the 8k certificate key using Importing and Replacing Certificates.
Client SSL Profile	Select the client-side SSL Profile from the drop-down menu.
SNI Certificates	Select the available SNI certificate from the drop-down menu.
Trusted CA Certificates	Select the available CA certificate.
Mandatory Client Authentication	Toggle the button to enable this menu item.
Certificate Chain Depth	Set the certificate chain depth to verify the depth in the server certificates chain.
Certificate Revocation List	Select the available CRL to disallow compromised server certificates.

6 Configure the Server SSL.

Option	Description
Server SSL	Toggle the button to enable the profile.
Client Certificate	Select a client certificate from the drop-down menu. This certificate is used if the server does not host multiple host names on the same IP address or if the client does not support Server Name Indication (SNI) extension.
Server SSL Profile	Select the Server-side SSL Profile from the drop-down menu.
Trusted CA Certificates	Select the available CA certificate.
Mandatory Server Authentication	Toggle the button to enable this menu item. Server-side SSL profile binding specifies whether the server certificate presented to the load balancer during the SSL handshake must be validated or not. When validation is enabled, the server certificate must be signed by one of the trusted CAs whose self-signed certificates are specified in the same server-side SSL profile binding.
Certificate Chain Depth	Set the certificate chain depth to verify the depth in the server certificates chain.
Certificate Revocation List	Select the available CRL to disallow compromised server certificates. OCSP and OCSP stapling are not supported on the server-side.

7 Click **Additional Properties** to configure additional Layer 7 virtual server properties.

Option	Description
Max Concurrent Connection	Set the maximum concurrent connection allowed to a virtual server so that the virtual server does not deplete resources of other applications hosted on the same load balancer.
Max New Connection Rate	Set the maximum new connection to a server pool member so that a virtual server does not deplete resources.
Sorry Server Pool	Select an existing sorry server pool from the drop-down menu. The sorry server pool serves the request when a load balancer cannot select a backend server to the serve the request from the default pool. You can click the vertical ellipses to create a server pool.
Default Pool Member Port	Enter a default pool member port, if the pool member port for a virtual server is not defined. For example, if a virtual server is defined with port range 2000-2999 and the default pool member port range is set as 8000-8999, then an incoming client connection to the virtual server port 2500 is sent to a pool member with a destination port set to 8500.
Admin State	Toggle the button to disable the admin state of the Layer 7 virtual server.
Access Log	Toggle the button to enable logging for the Layer 7 virtual server.
Log Significant Event Only	This field can only be configured if access logs are enabled. Requests with an HTTP response status of >=400 are treated as a significant event.
Tags	Select a tag from the drop-down list. You can specify a tag to set a scope of the tag.

8 Click **Save**.

Add Load Balancer Rules

With Layer 7 HTTP virtual servers, you can optionally configure load balancer rules and customize load balancing behavior using match or action rules.

Load balancer rules are supported for only Layer 7 virtual servers with an HTTP application profile. Different load balancer services can use load balancer rules.

Each load balancer rule consists of single or multiple match conditions and single or multiple actions. If the match conditions are not specified, then the load balancer rule always matches and is used to define default rules. If more than one match condition is specified, then the matching strategy determines if all conditions must match or any one condition must match for the load balancer rule to be considered a match.

Each load balancer rule is implemented at a specific phase of the load balancing processing; Transport, HTTP Access, Request Rewrite, Request Forwarding, and Response Rewrite. Not all the match conditions and actions are applicable to each phase.

Up to 4,000 load balancer rules can be configured with the API, if the `skip_scale_validation` flag in LbService is set. Note that the flag can be set via API. Refer to the *NSX API Guide* for more information. Up to 512 load balancer rules can be configured through the user interface.

Load Balancer rules support REGEX for match types. For more information, see Regular Expressions in Load Balancer Rules.

Prerequisites

Verify a Layer 7 HTTP virtual server is available. See Add Layer 7 HTTP Virtual Servers.

- Configure Transport Phase Load Balancer Rules

  Transport phase is the first phase of a client HTTP request.
- Configure HTTP Access Load Balancer Rules

  A JSON web token (JWT) is a standardized, optionally validated and/or encrypted format that is used to securely transfer information between two parties.
- Configure Request Rewrite Load Balancer Rules

  An HTTP request rewrite is applied to the HTTP request coming from the client.
- Configure Request Forwarding Load Balancer Rules

  Request forwarding redirects a URL or host to a specific server pool.
- Configure Response Rewrite Load Balancer Rules

  An HTTP response rewrite is applied to the HTTP response going out from the servers to the client.
- Regular Expressions in Load Balancer Rules

  Regular expressions (REGEX) are used in match conditions for load balancer rules.

Configure Transport Phase Load Balancer Rules

Transport phase is the first phase of a client HTTP request.

Load Balancer virtual server SSL configuration is found under **SSL Configuration**. There are two possible configurations. In both modes, the load balancer sees the traffic, and applies load balancer rules based on the client HTTP traffic.

- SSL Offload, configuring only the SSL client. In this mode, the client to VIP traffic is encrypted (HTTPS), and the load balancer decrypts it. The VIP to Pool member traffic is clear (HTTP).
- SSL End-to-End, configuring both the Client SSL and Server SSL. In this mode, the client to VIP traffic is encrypted (HTTPS), and the load balancer decrypts it and then re-encrypts it. The VIP to Pool member traffic is encrypted (HTTPS).

The Transport Phase is complete when the virtual server receives the client SSL hello message virtual server. this occurs before SSL is ended, and before HTTP traffic.

The Transport Phase allows administrators to select the SSL mode, annd specific server pool based on the client SSL hello message. There are three options for the virtual server SSL mode:

- SSL Offload
- End-to-End
- SSL-Passthrough (the load balancer does not end SSL)

Load Balancer rules support REGEX for match types. PCRE style REGEX patterns are supported with a few limitations on advanced use cases. When REGEX is used in match conditions, named capturing groups are supported. See Regular Expressions in Load Balancer Rules.

Prerequisites

Verify that a Layer 7 HTTP virtual server is available. See Add Layer 7 HTTP Virtual Servers.

Procedure

1. Open the Layer 7 HTTP virtual server.
2. In the Load Balancer Rules section, next to Transport Phase, click **Set > Add Rule** to configure the load balancer rules for the Transport Phase.
3. SSL SNI is the only match condition supported. Match conditions are used to match application traffic passing through load balancers.
4. From the drop-down list, select a **Match Type**: starts with, ends with, equals, contains, matches regex.
5. Enter a **SNI Name**.
6. Toggle the **Case Sensitive** button to set a case-sensitive flag for HTTP header value comparison.
7. Toggle the **Negate** button to enable it.
8. From the drop-down list, select a **Match Strategy**:

Match Strategy	Description
Any	Either host or path may match for this rule to be considered a match.
All	Both host and path must match for this rule to be considered a match.

9. From the drop-down menu, select the **SSL Mode Selection**.

SSL Mode	Description
SSL Passthrough	SSL Passthrough passes HTTP traffic to a backend server without decrypting the traffic on the load balancer. The data is kept encrypted as it travels through the load balancer. If SSL Passthrough is selected, a server pool can be selected. See Add a Server Pool for Load Balancing in Manager Mode.
SSL Offloading	SSL Offloading decrypts all HTTP traffic on the load balancer. SSL offloading allows data to be inspected as it passes between the load balancer and server. If NTLM and multiplexing are not configured, the load balancer establishes a new connection to the selected backend server for each HTTP request.
SSL End-to End	After receiving the HTTP request, the load balancer connects to the selected backend server and talks with it using HTTPS. If NTLM and multiplexing are not configured, the load balancer establishes a new connection to the selected backend server for each HTTP request.

10 Click **SAVE** and **APPLY**.

Configure HTTP Access Load Balancer Rules

A JSON web token (JWT) is a standardized, optionally validated and/or encrypted format that is used to securely transfer information between two parties.

In the HTTP ACCESS phase, users can define the action to validate JWT from clients and pass, or remove JWT to backend servers.

Load Balancer rules support REGEX for match types. PCRE style REGEX patters is supported with a few limitations on advanced use cases. When REGEX is used in match conditions, named capturing groups are supported. See Regular Expressions in Load Balancer Rules.

Prerequisites

Verify that a Layer 7 HTTP virtual server is available. See Add Layer 7 HTTP Virtual Servers.

Procedure

1 Open the Layer 7 HTTP virtual server.

2 In the Load Balancer Rules section, next to HTTP Access Phase, click **Set > Add Rule** to configure the load balancer rules for the HTTP Request Rewrite phase.

3 From the drop-down menu, select a match condition. Match conditions are used to match application traffic passing through load balancers. Multiple match conditions can be specified in one load balancer rule. Each match condition defines a criterion for application traffic.

Supported Match Condition	Description
HTTP Request Method	Match an HTTP request method. http_request.method - value to match
HTTP Request URI	Match an HTTP request URI without query arguments. http_request.uri - value to match
HTTP Request URI Arguments	Match an HTTP request URI query argument. http_request.uri_arguments - value to match
HTTP Request Version	Match an HTTP request version. http_request.version - value to match
HTTP Request Header	Match any HTTP request header. http_request.header_name - header name to match http_request.header_value - value to match
HTTP Request Cookie	Match any HTTP request cookie. http_request.cookie_value - value to match
HTTP Request Body	Match an HTTP request body content. http_request.body_value - value to match
TCP Header Port	Match a TCP source or the destination port. tcp_header.source_port - source port to match tcp_header.destination_port - destination port to match

Supported Match Condition	Description
IP Header Source	Matches IP header text boxes in of HTTP messages. The source type must be either a single IP address, a range of IP addresses, or a group. See Add a Group. ■ If IP Header Source is selected, with an IP Address source type, the source IP address of HTTP messages should match IP addresses which are configured in groups. Both IPv4 and IPv6 addresses are supported. ■ If IP Header Source is selected with a Group source type, select the group from the drop-down menu. ip_header.source_address - source address to match ip_header.destination_address - destination address to match
Variable	Create a variable and assign a value to the variable.
Client SSL	Match client SSL profile ID. ssl_profile_id - value to match
Case Sensitive	Set a case-sensitive flag for HTTP header value comparison. If true, case is significant when comparing HTTP body value.

4 From the drop-down list, select a **Match Type**: starts with, ends with, equals, contains, matches regex.

5 If needed, enter the URI.

6 From the drop-down list, select a **Match Strategy**:

Match Strategy	Description
Any	Either host or path may match for this rule to be considered a match.
All	Both host and path must match for this rule to be considered a match.

7 From the drop-down menu select an **Action**:

Action	Description
JWT Authentication	JSON Web Token (JWT) is an open standard that defines a compact and self-contained way for securely transmitting information between parties as a JSON object. This information can be verified and trusted because it is digitally signed. ■ Realm - A description of the protected area. If no realm is specified, clients often display a formatted hostname. The configured realm is returned when a client request is rejected with 401 http status. The response is: "WWW-Authentication: Bearer realm=<realm>". ■ Tokens - This parameter is optional. Load balancer searches for every specified token one-by-one for the JWT message until found. If not found, or if this text box is not configured, load balancer searches the Bearer header by default in the http request "Authorization: Bearer <token>" ■ Key Type - Symmetric key or asymmetric public key (certificate-id) ■ Preserve JWT - This is a flag to preserve JWT and pass it to backend server. If disabled, the JWT key to the backend server is removed.
Connection Drop	If negate is enabled, when Connection Drop is configured, all requests not matching the specified match condition are dropped. Requests matching the specified match condition are allowed.
Variable Assignment	Enables users to assign a value to a variable in HTTP Access Phase, in such a way that the result can be used as a condition in other load balancer rule phases.

8 Click **Save** and **Apply**.

Configure Request Rewrite Load Balancer Rules

An HTTP request rewrite is applied to the HTTP request coming from the client.

Prerequisites

Verify that a Layer 7 HTTP virtual server is available. See Add Layer 7 HTTP Virtual Servers.

Load Balancer rules support REGEX for match types. PCRE style REGEX patters is supported with a few limitations on advanced use cases. When REGEX is used in match conditions, named capturing groups are supported. See Regular Expressions in Load Balancer Rules.

Procedure

1 Open the Layer 7 HTTP virtual server.

2 In the Load Balancer Rules section, next to Request Rewrite Phase, click **Set > Add Rule** to configure the load balancer rules for the HTTP Request Rewrite phase.

3 From the drop-down list, select a match condition. Match conditions are used to match application traffic passing through load balancers. Multiple match conditions can be specified in one load balancer rule. Each match condition defines a criterion for application traffic.

Supported Match Condition	Description
HTTP Request Method	Match an HTTP request method. http_request.method - value to match
HTTP Request URI	Match an HTTP request URI without query arguments. http_request.uri - value to match
HTTP Request URI Arguments	Used to match URI arguments aka query string of HTTP request messages, for example, in URI http://exaple.com?foo=1&bar=2, the "foo=1&bar=2" is the query string containing URI arguments. In an URI scheme, query string is indicated by the first question mark ("?") character and terminated by a number sign ("#") character or by the end of the URI. http_request.uri_arguments - value to match
HTTP Request Version	Used to match the HTTP protocol version of the HTTP request messages http_request.version - value to match
HTTP Request Header	Used to match HTTP request messages by HTTP header fields. HTTP header fields are components of the header section of HTTP request and response messages. They define the operating parameters of an HTTP transaction. http_request.header_name - header name to match http_request.header_value - value to match
HTTP Request Cookie	Used to match HTTP request messages by cookie which is a specific type of HTTP header. The match_type and case_sensitive define how to compare cookie value. http_request.cookie_value - value to match
HTTP Request Body	Match an HTTP request body content. http_request.body_value - value to match
Client SSL	Match client SSL profile ID. ssl_profile_id - value to match
TCP Header Port	Match a TCP source or the destination port. tcp_header.source_port - source port to match tcp_header.destination_port - destination port to match
IP Header Source	Matches IP header fields in of HTTP messages. The source type must be either a single IP address, or a range of IP addresses, or a group. See Add a Group. ■ If IP Header Source is selected, with an IP Address source type, the source IP address of HTTP messages should match IP addresses which are configured in groups. Both IPv4 and IPv6 addresss are supported ■ If IP Header Source is selected with a Group source type, seelct the group from the drop-down list. ip_header.source_address - source address to match ip_header.destination_address - destination address to match

Supported Match Condition	Description
Variable	Create a variable and assign a value to the variable.
Case Sensitive	Set a case-sensitive flag for HTTP header value comparison. If true, case is significant when comparing HTTP body value.

4 From the drop-down menu, select a **Match Type**: starts with, ends with, equals, contains, or matches regex. Match type is used to match a condition with a specified action.

Match Type	Description
Starts With	If the match condition starts with the specified value, the condition matches.
Ends With	If the match condition ends with the specified value, the condition matches.
Equals	If the match condition is the same as the specified value, the condition matches.
Contains	If the match condition contains the specified value, the condition matches.
Matches Regex	If the match condition matches the specified values, the condition matches.

5 Specify the URI.

6 From the drop-down menu, select a **Match Strategy**:

Match Strategy	Description
Any	Indicates that either host or path can match for this rule to be considered a match.
All	Indicates that both host and path must match for this rule to be considered a match.

7 Select an Action from the drop-down menu:

Actions	Description
HTTP Request URI Rewrite	This action is used to rewrite URIs in matched HTTP request messages. Specify the URI and URI Arguments in this condition to rewrite the matched HTTP request message's URI and URI arguments to the new values. Full URI scheme of HTTP messages have following syntax: Scheme: [//[user[:password]@]host[:port]][/path][?query][#fragment The URI field of this action is used to rewrite the /path part in the above scheme. The URI Arguments field is used to rewrite the query part. Captured variables and built-in variables can be used in the URI and URI Arguments fields. a Enter the URI of the HTTP request b Enter the query string of URI, which typically contains key value pairs, for example: foo1=bar1&foo2=bar2.
HTTP Request Header Rewrite	This action is used to rewrite header fields of matched HTTP request messages to specified new values. a Enter the name of a header text box HTTP request message. b Enter the header value.

Actions	Description
HTTP Request Header Delete	This action is used to delete header fields of HTTP request messages at HTTP_REQUEST_REWRITE phase. One action can be used to delete all headers with same header name. To delete headers with different header names, multiple actions must be defined. ■ Enter the name of a header field of HTTP request message.
Variable Assignment	Create a variable and assign it a name and value.

8 Toggle the **Case Sensitive** button to set a case-sensitive flag for HTTP header value comparison.

9 Toggle the **Negate** button to enable it.

10 Click **Save** and **Apply**.

Configure Request Forwarding Load Balancer Rules

Request forwarding redirects a URL or host to a specific server pool.

Prerequisites

Verify that a Layer 7 HTTP virtual server is available. See Add Layer 7 HTTP Virtual Servers.

Load Balancer rules support REGEX for match types. PCRE style REGEX patters is supported with a few limitations on advanced use cases. When REGEX is used in match conditions, named capturing groups are supported. See Regular Expressions in Load Balancer Rules.

Procedure

1 Open the Layer 7 HTTP virtual server.

2 Click **Request Forwarding > Add Rule** to configure the load balancer rules for the HTTP Request Forwarding.

3 From the drop-down list, select a match condition. Match conditions are used to match application traffic passing through load balancers. Multiple match conditions can be specified in one load balancer rule. Each match condition defines a criterion for application traffic.

Supported Match Condition	Description
HTTP Request Method	Match an HTTP request method. http_request.method - value to match
HTTP Request URI	Match an HTTP request URI without query arguments. http_request.uri - value to match
HTTP Request URI Arguments	Used to match URI arguments aka query string of HTTP request messages, for example, in URI http://exaple.com?foo=1&bar=2, the "foo=1&bar=2" is the query string containing URI arguments. In an URI scheme, query string is indicated by the first question mark ("?") character and terminated by a number sign ("#") character or by the end of the URI. http_request.uri_arguments - value to match
HTTP Request Version	Used to match the HTTP protocol version of the HTTP request messages http_request.version - value to match

Supported Match Condition	Description
HTTP Request Header	Used to match HTTP request messages by HTTP header fields. HTTP header fields are components of the header section of HTTP request and response messages. They define the operating parameters of an HTTP transaction. http_request.header_name - header name to match http_request.header_value - value to match
HTTP Request Cookie	Used to match HTTP request messages by cookie which is a specific type of HTTP header. The match_type and case_sensitive define how to compare cookie value. http_request.cookie_value - value to match
HTTP Request Body	Match an HTTP request body content. http_request.body_value - value to match
Client SSL	Match client SSL profile ID. ssl_profile_id - value to match
TCP Header Port	Match a TCP source or the destination port. tcp_header.source_port - source port to match tcp_header.destination_port - destination port to match
IP Header Source	Matches IP header fields in of HTTP messages. The source type must be either a single IP address, or a range of IP addresses, or a group. See Add a Group. ■ If IP Header Source is selected, with an IP Address source type, the source IP address of HTTP messages should match IP addresses which are configured in groups. Both IPv4 and IPv6 addresss are supported ■ If IP Header Source is selected with a Group source type, seelct the group from the drop-down list. ip_header.source_address - source address to match ip_header.destination_address - destination address to match
Variable	Create a variable and assign a value to the variable.
Case Sensitive	Set a case-sensitive flag for HTTP header value comparison. If true, case is significant when comparing HTTP body value.

4 Select an action:

Action	Description
HTTP Reject	Used to reject HTTP request messages. The specified reply_status value is used as the status code for the corresponding HTTP response message. The response message is sent back to client (usually a browser) indicating the reason it was rejected. http_forward.reply_status - HTTP status code used to reject http_forward.reply_message - HTTP rejection message
HTTP Redirect	Used to redirect HTTP request messages to a new URL. The HTTP status code for redirection is 3xx, for example, 301, 302, 303, 307, etc. The redirect_url is the new URL that the HTTP request message is redirected to. http_forward.redirect_status - HTTP status code for redirect http_forward.redirect_url - HTTP redirect URL

Action	Description
Select Pool	Force the request to a specific server pool. Specified pool member's configured algorithm (predictor) is used to select a server within the server pool. The matched HTTP request messages are forwarded to the specified pool. When HTTP keep-alive is enabled and forwarding rules are configured in the load balancer, the server keep-alive setting takes precedence. As a result, HTTP requests are sent to servers already connected with keep-alive. If you always want to give priority to the forwarding rules when the load balancer rule conditions are met, disable the keep-alive setting. Note that the persistence setting takes precedence over the keep-alive setting. Processing is done in the order of Persistence > Keep-Alive > Load Balancer Rules http_forward.select_pool - server pool UUID
Variable Persistence On	Select a generic persistence profile and enter a variable name. You can also enable **Hash Variable**. If the variable value is long, hashing the variable ensures that it is correctly stored in the persistence table. If the **Hash Variable** is not enabled, only the fixed prefix part of the variable value is stored in the persistence table if the variable value is long. As a result, two different requests with long variable values might be dispatched to the same backend server because their variable values have the same prefix part, when they should be dispatched to different backend servers.
Connection Drop	If negate is enabled in condition, when Connection Drop is configured, all requests not matching the condition are dropped. Requests matching the condition are allowed.
Reply Status	Shows the status of the reply.
Reply Message	Server responds with a reply message that contains confirmed addresses and configuration.

5 Click **Save** and **Apply**.

Configure Response Rewrite Load Balancer Rules

An HTTP response rewrite is applied to the HTTP response going out from the servers to the client.

Prerequisites

Verify that a Layer 7 HTTP virtual server is available. See Add Layer 7 HTTP Virtual Servers.

Load Balancer rules support REGEX for match types. PCRE style REGEX patters is supported with a few limitations on advanced use cases. When REGEX is used in match conditions, named capturing groups are supported. See Regular Expressions in Load Balancer Rules.

Procedure

1 Open the Layer 7 HTTP virtual server.

2 Click **Response Rewrite > Add Rule** to configure the load balancer rules for the HTTP Response Rewrite.

All match values accept regular expressions.

Supported Match Condition	Description
HTTP Response Header	This condition is used to match HTTP response messages from backend servers by HTTP header fields. http_response.header_name - header name to match http_response.header_value - value to match
HTTP Response Method	Match an HTTP response method. http_response.method - value to match
HTTP Response URI	Match an HTTP response URI. http_response.uri - value to match
HTTP Response URI Arguments	Match an HTTP response URI arguments. http_response.uri_args - value to match
HTTP Response Version	Match an HTTP response version. http_response.version - value to match
HTTP Response Cookie	Match any HTTP response cookie. http_response.cookie_value - value to match
Client SSL	Match client SSL profile ID. ssl_profile_id - value to match
TCP Header Port	Match a TCP source or the destination port. tcp_header.source_port - source port to match tcp_header.destination_port - destination port to match
IP Header Source	Matches IP header fields in of HTTP messages. The source type must be either a single IP address, or a range of IP addresses, or a group. See Add a Group. The source IP address of HTTP messages should match IP addresses which are configured in groups. Both IPv4 and IPv6 addresses are supported. ip_header.source_address - source address to match ip_header.destination_address - destination address to match
Variable	Create a variable and assign a value to the variable.
Case Sensitive	Set a case-sensitive flag for HTTP header value comparison.

3 Select an action:

Action	Description
HTTP Response Header Rewrite	This action is used to rewrite header fields of HTTP response messages to specified new values. http_response.header_name - header name http_response.header_value - value to write
HTTP Response Header Delete	This action is used to delete header fields of HTTP response messages. http_request.header_delete - header name http_request.header_delete - value to write
Variable Persistence Learn	Select a generic persistence profile and enter a variable name. You can also enable **Hash Variable**. If the variable value is long, hashing the variable ensures that it will be correctly stored in the persistence table. If **Hash Variable** is not enabled, only the fixed prefix part of the variable value is stored in the persistence table if the variable value is long. As a result, two different requests with long variable values might be dispatched to the same backend server (because their variable values have the same prefix part) when they should be dispatched to different backend servers.

4 Click **Save** and **Apply**.

Regular Expressions in Load Balancer Rules

Regular expressions (REGEX) are used in match conditions for load balancer rules.

Perl Compatible Regular Expressions (PCRE) style REGEX patterns is supported with a few limitations on advanced use cases. When REGEX is used in match conditions, named capturing groups are supported.

REGEX restrictions include:

- Character unions and intersections are not supported. For example, do not use [a-z[0-9]] and [a-z&&[aeiou]] instead use [a-z0-9] and [aeiou] respectively.
- Only 9 back references are supported and \1 through \9 can be used to refer to them.
- Use \0dd format to match octal characters, not the \ddd format.
- Embedded flags are not supported at the top level, they are only supported within groups. For example, do not use "Case (?i:s)ensitive" instead use "Case ((?i:s)ensitive)".
- Preprocessing operations \l, \u, \L, \U are not supported. Where \l - lowercase next char \u - uppercase next char \L - lower case until \E \U - upper case to \E.
- (?(condition)X), (?{code}), (??{Code}) and (?#comment) are not supported.
- Predefined Unicode character class \X is not supported
- Using named character construct for Unicode characters is not supported. For example, do not use \N{name} instead use \u2018.

  When REGEX is used in match conditions, named capturing groups are supported. For example, REGEX match pattern /news/(?<year>\d+)-(?<month>\d+)-(?<day>\d+)/(?<article>.*) can be used to match a URI like /news/2018-06-15/news1234.html.

Then variables are set as follows, $year = "2018" $month = "06" $day = "15" $article = "news1234.html". After the variables are set, these variables can be used in load balancer rule actions. For example, URI can be rewritten using the matched variables like, /news.py?year=$year&month=$month&day=$day&article=$article. Then the URI gets rewritten as /news.py?year=2018&month=06&day=15&article=news1234.html.

Rewrite actions can use a combination of named capturing groups and built-in variables. For example, URI can be written as /news.py?year=$year&month=$month&day=$day&article=$article&user_ip=$_remote_addr. Then the example URI gets rewritten as /news.py?year=2018&month=06&day=15&article=news1234.html&user_ip=1.1.1.1.

Note For named capturing groups, the name cannot start with an _ character.

In addition to named capturing groups, the following built-in variables can be used in rewrite actions. All the built-in variable names start with _.

- $_args - arguments from the request
- $_arg_<name> - argument <name> in the request line
- $_cookie_<name> - value of <name> cookie
- $_upstream_cookie_<name> - cookie with the specified name sent by the upstream server in the "Set-Cookie" response header field
- $_upstream_http_<name> - arbitrary response header field and <name> is the field name converted to lower case with dashes replaced by underscores
- $_host - in the order of precedence - host name from the request line, or host name from the "Host" request header field, or the server name matching a request
- $_http_<name> - arbitrary request header field and <name> is the field name converted to lower case with dashes replaced by underscores
- $_https - "on" if connection operates in SSL mode, or "" otherwise
- $_is_args - "?" if a request line has arguments, or "" otherwise
- $_query_string - same as $_args
- $_remote_addr - client address
- $_remote_port - client port
- $_request_uri - full original request URI (with arguments)
- $_scheme - request scheme, "http" or "https"
- $_server_addr - address of the server which accepted a request
- $_server_name - name of the server which accepted a request
- $_server_port - port of the server which accepted a request
- $_server_protocol - request protocol, usually "HTTP/1.0" or "HTTP/1.1"

- $_ssl_client_escaped_cert - returns the client certificate in the PEM format for an established SSL connection.
- $_ssl_server_name - returns the server name requested through SNI
- $_uri - URI path in request
- $_ssl_ciphers: returns the client SSL ciphers
- $_ssl_client_i_dn: returns the "issuer DN" string of the client certificate for an established SSL connection according to RFC 2253
- $_ssl_client_s_dn: returns the "subject DN" string of the client certificate for an established SSL connection according to RFC 2253
- $_ssl_protocol: returns the protocol of an established SSL connection
- $_ssl_session_reused: returns "r" if an SSL session was reused, or "." otherwise

Groups Created for Server Pools and Virtual Servers

NSX Manager automatically creates groups for load balancer server pools and VIP ports.

Load Balancer created groups are visible under **Inventory > Groups**.

Server pool groups are created with the name NLB.PoolLB.*Pool_Name LB_Name* with group member IP addresses assigned:

- Pool configured with no LB-SNAT (transparent): 0.0.0.0/0
- Pool configured with LB-SNAT Automap: T1-Uplink IP 100.64.x.y and T1-ServiceInterface IP
- Pool configured with LB-SNAT IP-Pool: LB-SNAT IP-Pool

VIP Groups are created with the name NLB.VIP.*virtual server name* and the VIP group member IP addresses are VIP IP@.

For server pool groups, you can create an allow traffic distributed firewall rule from the load balancer (NLB.PoolLB. *Pool_Name LB_Name*). For Tier-1 gateway firewall, you can create an allow traffic from clients to LB VIP NLB.VIP.*virtual server name*.

Distributed Load Balancer

11

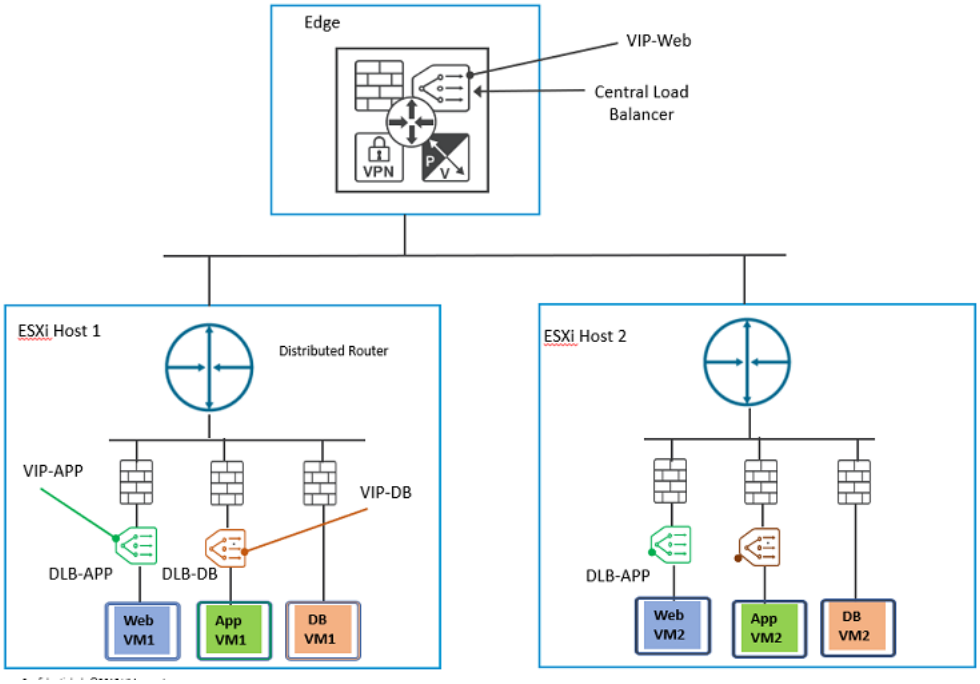
A Distributed Load Balancer configured in NSX can help you effectively load balance East-West traffic and scale traffic because it runs on each ESXi host.

Important Distributed Load Balancer is supported only for Kubernetes (K8s) cluster IPs managed by vSphere with Kubernetes. Distributed Load Balancer is not supported for any other workload types. As an administrator, you cannot use NSX Manager GUI to create or modify Distributed Load Balancer objects. These objects are pushed by VMware vCenter through NSX API when K8 cluster IPs are created in VMware vCenter.

In traditional networks, a central load balancer deployed on an NSX Edge node is configured to distribute traffic load managed by virtual servers that are configured on the load balancer.

If you are using a central balancer, increasing the number of virtual servers in the load balancer pool might not always meet scale or performance criteria for a multi-tier distributed application. A distributed load balancer is realized on each hypervisor where load balancing workloads, such as clients and servers are deployed, ensuring traffic is load balanced on each hypervisor in a distributed way.

A distributed load balancer can be configured on the NSX network along with a central load balancer.



In the diagram, an instance of the Distributed Load Balancer is attached to a VM group. As the VMs are downlinks to the distributed logical router, Distributed Load Balancer only load balances east-west traffic. In contrast, the central load balancer, manages north-south traffic.

To cater load balancing requirements of each component or module of an application, a distributed load balancer can be attached to each tier of an application. For example, to serve a user request, a frontend of the application needs to reach out to the middle module to get data. However, the middle layer might not be deployed to serve the final data to the user, so it needs to reach out the backend layer to get additional data. For a complex application, many modules might need to interact with each other to get information. Along with complexity, when the number of user request increase exponentially, a distributed load balancer can efficiently meet the user needs without taking a performance hit. Configuring a Distributed Load Balancer on every host achieves issues of scale and packet transmission efficiency.

Important Enable DFW for DLB workloads. Disabling DFW either globally or through the DFW Exclusion List will cause an outage on DLB workloads.

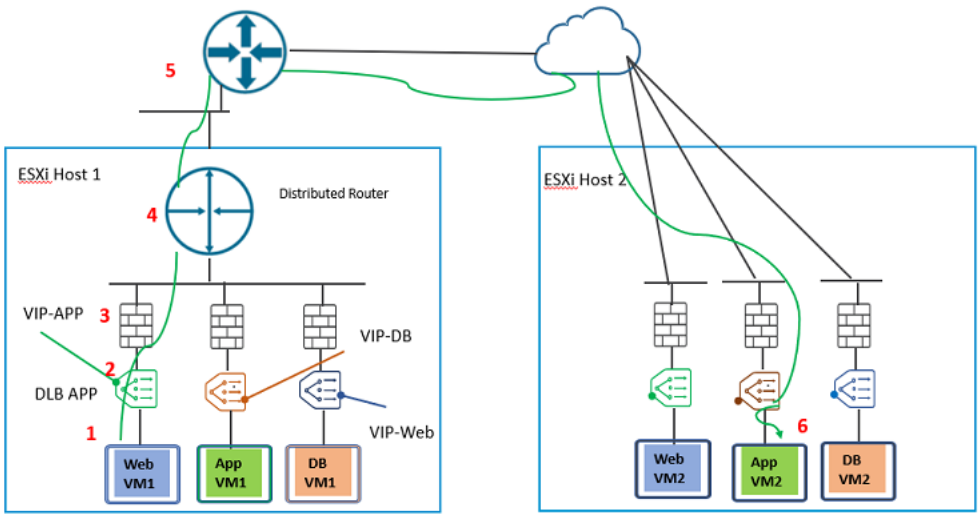
This chapter includes the following topics:

- Understanding Traffic Flow with a Distributed Load Balancer
- Create and Attach a Distributed Load Balancer Instance
- Create a Server Pool for Distributed Load Balancer
- Create a Virtual Server with a Fast TCP or UDP Profile
- Verifying Distributed Load Balancer Configuration on ESXi Hosts
- Distributed Load Balancer Statistics and Diagnostics

- Distributed Load Balancer Operational Status
- Run Traceflow on Distributed Load Balancer
- Supported Features

Understanding Traffic Flow with a Distributed Load Balancer

Understand how traffic flows between VMs that are connected to an instance of a distributed load balancer (DLB).



As an administrator ensure:

- Virtual IP addresses and pool members connected to a DLB instance must have unique IP address for traffic to be routed correctly.

Traffic flow between Web VM1 and APP VM2.

1 When Web VM1 sends out a packet to APP VM2 it is received by the VIP-APP.

   The DLB APP is attached to the policy group consisting of Web tier VMs. Similarly, DLB-APP hosting VIP-DB must be attached to the policy group consisting of App tier VMs.

2 The VIP-APP hosted on DLB APP receives the request from Web VM1.

3 Before reaching the destination VM group, the packet is filtered by distributed firewall rules.

4 After the packets are filtered based on the firewall rules, it is sent to the Tier-1 router.

5 It is further routed to the the physical router.

6 The route is completed when the packet is delivered to the destination App VM2 group.

As DLB VIPs can only be accessed from VMs connected to downlinks of Tier-0 or Tier-1 logical routers, DLB provides load balancing services to east-west traffic.

A DLB instance can co-exist with an instance of DFW. With DLB and DFW enabled on a virtual interface of a hypervisor, first the traffic is load balanced based on the configuraiton in DLB and then DFW rules are applied on traffic flowing from a VM to the hypervisor. DLB rules are applied on traffic originating from downlinks of a Tier-0 or Tier-1 logical routers going to the destination hypervisor. DLB rules cannot be applied on traffic flowing in the reverse direction - originating from outside the host going to a destination VM.

For example, if the DLB instance is load balancing traffic from Web-VMs to App-VMs, then to allow such traffic to pass through DFW, ensure that the DFW rule is set to value "Source=Web-VMs, Destination=App-VMs, Action=Allow".

Create and Attach a Distributed Load Balancer Instance

Unlike a central load balancer, a Distributed Load Balancer (DLB) instance is attached to virtual interfaces of a VM group.

At the end of the procedure a DLB instance is attached to the virtual interfaces of a VM group. It is only possible to create and attach a DLB instance through API commands.

Prerequisites

- Add a policy group consisting of VMs. For example, such a VM group can be related to the App tier that receives requests from a VM on the Web-tier.

Procedure

- Run `Put /policy/api/v1/infra/lb-services/<mydlb>`.

  {

  "connectivity_path" : "/infra/domains/default/groups/<clientVMGroup>",

  "enabled" : true,

  "size" : "DLB",

  "error_log_level" : "Debug",

  "access_log_enabled" : false,

  "resource_type" : "LBService",

  "display_name" : "mydlb"

  }

  Where,

  - `connectivity_path`:
    - If the connectivity path is set to **`Null`** or **`Empty`**, the DLB instance is not applied to any transport nodes.

    - If the connectivity path is set **ALL**, all virtual interfaces of all transport nodes are bound to the DLB instance. One DLB instance is applied to all the virtual interfaces of the policy group.
- `size`: Set to value **DLB**. As each application or virtual interface gets an instance of DLB, there is just a single size form factor of the DLB instance.
- `enabled`: By default, the created DLB instance is enabled. You cannot disable the DLB instance.
- `error_log_level`: Supported levels are **Debug**, **Error**, and **Info**. By default, log level is set to **Info**. To get verbose logs, set the level to **Debug**.

A DLB instance is created and attached to the VM group. The DLB instance created on the Web-tier is attached to all the virtual interfaces of the Web-tier VM group.

What to do next

After creating a DLB instance, log in to the NSX Manager, go to **Networking -> Load Balancing -> Load Balancers**. View details of the DLB instance.

Next, Create a Server Pool for Distributed Load Balancer.

Create a Server Pool for Distributed Load Balancer

Create a load balancer pool to include virtual machines that consume DLB services.

This task can be done both from the NSX UI and NSX API.

The API command to create a DLB pool is `PUT https://<NSXManager_IPaddress>/policy/api/v1/infra/lb-pools/<lb-pool-id>`

Prerequisites

- Create a VM group that consumes DLB service.
- Create and attach a DLB instance to a VM group.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Go to **Networking > Load Balancing > Server Pools**.
3. Click **Add Server Pool**.

4 Enter values in these fields.

Field	Description
Name	Enter name of the DLB pool.
Algorithm	**Weighted Round Robin**, **Round Robin**, **Weighted Least Connection**, **Least Connection** and **IP Hash** are the supported algorithms. Since Distributed Load Balancer runs locally on each ESXi server, these algorithms are local to each ESXi server. There is no synchronization of load balancing connection information between different ESXi servers of a cluster. ■ **Weighted Round Robin**: Use this algorithm to send connections to pool members based on the weights assigned to each pool member. For example, if you assign pool member A with weight 3, pool member B with weight 2 and pool member C with weight 1, then out of a total of 6 client connections, pool member A receives 3 connections, pool member B receives 2 connections and pool member C receives 1 connection. ■ **Round Robin**: Use this algorithm to send equal number of connections to each pool member. ■ **Least Connection**: Use this algorithm so that a pool member with the least number of active connections. Each pool member is configured to a slow start (Slow Start is set to True). When it receives connections, the status of the pool member is set to Slow Start is False. ■ **Weighted Least Connection**: Use this algorithm, to send connections to pool members based on the weights assigned to each pool member. ■ **IP Hash**: Use this algorithm to send connections based on the hash of IP addresses. **Note** Do not use **IP Hash** if you want to persist connections to the same pool member even after the number of pool members change.
Members/Group	Click **Select Members** and on the **Configure Server Pool Members** window, do one of the following: ■ Select **Enter individual members**. To add a new member, click **Add Member** and enter values in the mandatory fields. ■ Select **Select a group** and **Add Group** or select an existing group. To add a new group, enter values in these fields. ■ **Name** ■ **Compute Members**: Click **Set Members** to add a group that includes all the pool members. ■ **IP Revision Filter**: Both IPv4 and IPv6 are supported. ■ **Port**: Default port for all the dynamic pool members.
SNAT Translation Mode	Set this field to **Disabled** state. SNAT translation is not supported in a Distributed Load Balancer.

5 Click **Save**.

Results

Server pool members are added for the Distributed Load Balancer.

What to do next

See Create a Virtual Server with a Fast TCP or UDP Profile.

Create a Virtual Server with a Fast TCP or UDP Profile

Create a virtual server and bind it to a Distributed Load Balancer service.

This task can be performed both from the NSX UI and NSX APIs.

The API command to create a virtual server is `PUT https://<NSXManager_IPAddress>/policy/api/v1/infra/lb-virtual-servers/<lb-virtual-server-id>`.

Prerequisites

- Create a server pool for the Distributed Load Balancer.
- To use IPv6 addresses as the virtual IP of Distributed Load Balancer, on the Global Networking Config page (**Networking > Global Networking Config**), ensure **L3 Forwarding Mode** to **IPv4 and IPv6**.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Go to **Networking** → **Load Balancing** → **Virtual Servers**.
3. Click **Add Virtual Server -> L4 TCP**.
4. To configure a virtual server for a Distributed Load Balancer, only the following fields are supported.

Field	Description
Name	Enter a name for the virtual server.
IP Address	Supports both IPv4 and IPv6 addresses. Enter the IP address of the Distributed Load Balancer virtual server. All client connections arrive at this IP address of the Distributed Load Balancer virtual server.
Ports	Virtual server port number. Multiple ports or port ranges are not supported in the virtual server of a Distributed Load Balancer.
Load Balancer	Attach the Distributed Load Balancer instance that is associated to the virtual server. The virtual server then knows which policy group the load balancer is servicing.
Server Pool	Select the server pool. The server pool contains backend servers. Server pool consists of one or more servers that are similarly configured and are running the same application. It is also referred to as pool members. **Note** If the virtual IP address of the Distributed Load Balancer is IPv4, the server pool members must be of the same versions. Likewise if you use IPv6 version of virtual IP address.

Field	Description
Application Profile	Select the application profile for the virtual server. The application profile defines the application protocol characteristics. It is used to influence how load balancing is performed. The supported application profiles are: ■ Load Balancer Fast TCP Profile ■ Load Balancer Fast UDP Profile
Default Pool Member Ports	Optional field. Enter one port number to be used when member ports are not defined. Multiple ports or port ranges for default pool member ports are not supported in the virtual server of a Distributed Load Balancer.
Persistence	Optional field. Select **Source IP** or **Disabled**.

The Distributed Load Balancer configuration is complete.

Results

Verify whether the DLB is distributing traffic to all the servers in the pool based on the algorithm defined in the configuration. If you choose the Round_Robin algorithm, then DLB must be able to choose servers from the pool in a round robin fashion.

In the ESXi host, verify whether the DLB configuration is complete.

What to do next

See Verifying Distributed Load Balancer Configuration on ESXi Hosts.

Verifying Distributed Load Balancer Configuration on ESXi Hosts

Verify whether the Distributed Load Balancer was configured completely on ESXi hosts.

After you securely connect to the ESXi host, run `/opt/vmware/nsx-nestdb/bin/nestdb-cli`. From the `nestdb-cli` prompt, run the following commands.

Command	Sample Response
To view the configured DLB service, run `get LbServiceMsg`.	`{'id': {'left': 13946864992859343551, 'right': 10845263561610880178}, 'virtual_server_id': [{'left': 13384746951958284821, 'right': 11316502527836868364}], 'display_name': 'mydlb', 'size': 'DLB', 'enabled': True, 'access_log_enabled': False, 'log_level': 'LB_LOG_LEVEL_INFO', 'applied_to': {'type': 'CONTAINER', 'attachment_id': {'left': 2826732686997341216, 'right': 10792930437485655035}}}`
To view the virtual server configured for DLB, run `get LbVirtualServerMsg`.	`{'port': '80', 'revision': 0, 'display_name': 'mytcpvip', 'pool_id': {'left': 4370937730160476541, 'right': 13181758910457427118}, 'enabled': True, 'access_log_enabled': False, 'id': {'left': 13384746951958284821, 'right': 11316502527836868364}, 'ip_protocol': 'TCP', 'ip_address': {'ipv4': 2071690107}, 'application_profile_id': {'left': 1527034089224553657, 'right': 10785436903467108397}}`
To view configuration of the DLB pool members, run `get LbPoolMsg`.	`{'tcp_multiplexing_number': 6, 'display_name': 'mylbpool', 'tcp_multiplexing_enabled': False, 'member': [{'port': '80', 'weight': 1, 'display_name': 'Member_VM30', 'admin_state': 'ENABLED', 'ip_address': {'ipv4': 3232261280}, 'backup_member': False}, {'port': '80', 'weight': 1, 'display_name': 'Member_VM31', 'admin_state': 'ENABLED', 'ip_address': {'ipv4': 3232261281}, 'backup_member': False}, {'port': '80', 'weight': 1, 'display_name': 'Member_VM32', 'admin_state': 'ENABLED', 'ip_address': {'ipv4': 3232261282}, 'backup_member': False}], 'id': {'left': 4370937730160476541, 'right': 13181758910457427118}, 'min_active_members': 1, 'algorithm': 'ROUND_ROBIN'}`

Command	Sample Response
To view NSX controller configuration pushed to the ESXi host, run `get ContainerMsg`.	`{'container_type': 'CONTAINER', 'id': {'left': 2826732686997341216, 'right': 10792930437485655035}, 'vif': ['cd2e482b-2998-480f-beba-65fbd7ab1e62', 'f8aa2a58-5662-4c6b-8090-d1bd19174205', '83a1f709-e675-4e42-b677-ff501fd0f4ec', 'b8366b39-4c81-41fc-b89e-de7716462b2f'], 'name': 'default.clientVMGroup', 'mac_address': [{'mac': 52237218275}, {'mac': 52243694681}, {'mac': 52233233291}, {'mac': 52239463383}], 'ip_address': [{'ipv4': 16844388}, {'ipv4': 16844644}, {'ipv4': 16844132}, {'ipv4': 3232261283}, {'ipv4': 16844298}, {'ipv4': 16844554}, {'ipv4': 16844042}]}`
To view application profile configuration on the ESXi host, run `get LbApplicationProfileMsg`.	`{'display_name': 'default-tcp-lb-app-profile', 'id': {'left': 1527034089224553657, 'right': 10785436903467108397}, 'application_type': 'FAST_TCP', 'fast_tcp_profile': {'close_timeout': 8, 'flow_mirroring_enabled': False, 'idle_timeout': 1800}}`

Distributed Load Balancer Statistics and Diagnostics

NSX API and CLI commands to monitor statistics for Distributed Load Balancer instances.

CLI Commands for Distributed Load Balancer

Action	Command
Display all load balancers.	`get load-balancers`
Display a specific load balancer.	`get load-balancer <UUID_LoadBalancer>`
Show statistics of all pools of the specified load balancer	`get load-balancer <UUID_LoadBalancer> pools stats`
Show load balancer virtual-server configuration.	`get load-balancer <UUID_LoadBalancer> virtual-servers`
Show statistics of the specified load balancer and pool	`get load-balancer <UUID_LoadBalancer> pool <UUID_Pool> stats`
Show persistence-tables entry	`get load-balancer <UUID_LoadBalancer> persistence-tables`
Show load balancer pools configuration	`get load-balancer <UUID_LoadBalancer> pools`
Show statistics of all virtual servers of the specified load balancer	`get load-balancer <UUID_LoadBalancer> virtual-servers stats`
Show statistics of the specified load balancer and virtual server	`get load-balancer <UUID_LoadBalancer> virtual-server <UUID_VirtualSerever> stat`

Action	Command
Clear statistics of the specified load balancer and pool	`clear load-balancer <UUID_LoadBalancer> pool <UUID_Pool> stats`
Clear statistics of all pools of the specified load balancer	`clear load-balancer <UUID_LoadBalancer> pools stats`
Clear statistics of the specified load balancer	`clear load-balancer <UUID_LoadBalancer> stats`
Clear statistics of the specified load balancer and virtual server	`clear load-balancer <UUID_LoadBalancer> virtual-server <UUID_VirtualServer> stats`
Clear statistics of all virtual servers of the specified load balancer	`clear load-balancer <UUID_LoadBalancer> virtual-servers stats`
View L4 session table details	`get load-balancer <UUID_LoadBalancer> sessions`
Display distributed load balancer statistics	`get load-balancer <UUID_LoadBalancer> stats`

CLI Diagnostic Commands for Distributed Load Balancer

Action	Command
Show load balancer diagnosis information	`get load-balancer <UUID_LoadBalancer> diagnosis`

This command runs a diagnosis report on the following:

- Checking System:
  - Edge memory usage. If edge memory usage is higher than 90%, memory usage is shown in the result.
  - Disk usage. Only "/", "/var/log" and "/config" folders are checked. If disk usage for any folder is higher than 90%, disk usage of this folder is shown in the result.
- Checking Crash - if a core file is generated by datapathd, LB nginx, lb-dispatcher, lb_conf, nsx-edge-exporter, nsd or lbconf_gen, the core file name is listed in the result.
- Checking Daemon Status - process dispatcher, datapathd, nsxa, nsd, nestdb, and LB nginx are checked. If any of these processes are not running, it is listed in the result.
- Checking Configuration:
  - Whether this LB object has been in nestdb.
  - Whether there is FATAL cfg information related this LB instance in syslog.
  - Whether there is configuration build failure in lbconf_gen.log.
  - Whether there is LbCurrentMsg object in nestdb.
  - Whether the generation_id in nginx.conf is equal with the value in LbCurrentMsg.
  - Whether LB firewall rules have been written into nestdb.
  - Whether LB firewall rules are working in datapathd.

- Checking runtime - health check status of pool member is checked. If the status of the health check is down, it is reported in the result.
- Checking Stats - datapath status, kni stats, and LB nginx stats are checked.
  - If there is an error in the LB cache from datapathd, it is reported in the result.
  - If there is an rx_drop or tx_drop in the LB kni interface, it is listed in the result.
  - If there is an error in the LB nginx statistics, it is reported in the result.

API commands for Distributed Load Balancer

Action	Command
Get Distributed Load Balancer statistics **Note** If you do not specify an enforcement point path, API fetches information from all enforcement paths and displays aggregated information. If you specify an enforcement point path, only information for that path is retrieved and displayed in the response body of the API call.	`GET /policy/api/v1/infra/lb-services/LB_Service/statistics?source=realtime&enforcement_point_path=/infra/sites/default/enforcement-points/default`
Get Distributed Load Balancer virtual server statistics	`GET /infra/lb-services/<lb-service-id>/lb-virtual-servers/<lb-virtual-server-id>/statistics?source=realtime&enforcement_point_path=/infra/sites/default/enforcement-points/default`
Get Distributed Load Balancer pool statistics	`GET /infra/lb-services/<lb-service-id>/lb-pools/<lb-pool-id>/statistics?source=realtime&enforcement_point_path=/infra/sites/default/enforcement-points/default`

Distributed Load Balancer Operational Status

Know the operational status of the distributed load balancer service in NSX Manager UI and on ESXi hosts.

As distributed load balancer service scales linearly as the number of hosts increases ESXi, a single distributed load balancer service can support several ESXi hosts. In turn, each ESXi host can support multiple virtual interfaces (VIFs), across many ESXi hosts. The consolidated status of distributed load balancer at NSX Manager level is calculated using the consolidated status of all the associated ESXi hosts. The consolidated status of distributed load balancer at ESXi host level is calculated using the individual status of all associated VIFs on that ESXi host.

Status at NSX Manager

Status	Description
`Up`	Status is `Up` when all the related transport nodes return status for the distributed load balancer service as `ready`.
`Degraded`	Status is `Degraded` when all the following conditions are true: ■ At least one transport node returns status for the distributed load balancer service as `ready` or `partially ready` ■ Not all the related transport nodes return status for the load balancer service as `ready`.
`Down`	Status is `Down` when one of the following conditions is true: ■ All the related transport nodes return `not ready`. ■ At least one transport node returns `not ready` and no transport node returns `ready`.
`Unknown`	Status is `Unknown` when all the related transport nodes return status for the distributed load balancer service as `Unknown`.
`Disabled`	Status is `Disabled` when the distributed load balancer service is enabled but the connectivity path is not specified.

Status at ESXi Host

Status	Description
`ready`	The consolidated status for the distributed load balancer service on the ESXi Host is `ready` when the status of all associated VIFs on this ESXi Host are `ready`. **Note** ■ `ready` status on VIF means that the distributed load balancer instance is the oldest and applied.
`not ready`	The consolidated status for the distributed load balancer service on the ESXi Host is `not ready` when no associated VIF is `ready`.
`partially ready`	The consolidated status for the distributed load balancer service on the ESXi Host is `partially ready` when both of the following conditions are true: ■ At least one associated VIF is `ready`. ■ At least one associated VIF is `not ready` or `conflict` . **Note** ■ `not ready` status on VIF means that the distributed load balancer service instance is the oldest, should be applied, but not applied. ■ `conflict` status on VIF means that the distributed load balancer service instance is not the oldest and not applied.

Detailed Status Through API

Run the following API to get detailed status of distributed load balancer instance running at a transport node.

```
GET https://
<manager IP>/policy/api/v1/infra/lb-services/<DLBname>/detailed-status?
source=realtime&include_instance_details=true&transport_node_ids=node1_uuid
```

Sampled response:

```
{
  "results":
    {
      "service_path": "/infra/lb-services/mydlb",
      "service_status": "UP",
      "virtual_servers": [
        {
          "virtual_server_path": "/infra/lb-virtual-servers/mytcpvip",
          "status": "UP",
          "last_update_timestamp": 1591344963509,
          "resource_type": "LBVirtualServerStatus"
        }
      ],
      "pools": [
        {
          "pool_path": "/infra/lb-pools/mylbpool",
          "status": "UP",
          "last_update_timestamp": 1591344963509,
          "resource_type": "LBPoolStatus"
        }
      ],
      "last_update_timestamp": 1591344963509,
      "instance_detail_per_tn": [
        {
          "transport_node_id": "b09b7b6c-a60d-11ea-835e-d95476fe6438",
          "instance_detail_per_status": [
            {
              "status": "READY",
              "instance_number": 3,
              "instance_details": [
                {
                  "attachment_display_name": "12-vm_Client_VM_Ubuntu_1404-local-1762/12-
vm_Client_VM_Ubuntu_1404-local-1762.vm@b09b7b6c-a60d-11ea-835e-d95476fe6438"
                },
                {
                  "attachment_display_name": "10-vm_Client_VM_Ubuntu_1404-local-1762/10-
vm_Client_VM_Ubuntu_1404-local-1762.vm@b09b7b6c-a60d-11ea-835e-d95476fe6438"
                },
                {
                  "attachment_display_name": "11-vm_Client_VM_Ubuntu_1404-local-1762/11-
vm_Client_VM_Ubuntu_1404-local-1762.vm@b09b7b6c-a60d-11ea-835e-d95476fe6438"
                }
              ]
            },
```

```
            {
              "status": "NOT_READY",
              "instance_number": 0
            },
            {
              "status": "CONFLICT",
              "instance_number": 0
            }
          ]
        }
      ],
      "enforcement_point_path": "/infra/sites/default/enforcement-points/default",
      "resource_type": "LBServiceStatus"
    }
  ],
  "intent_path": "/infra/lb-services/mydlb"
}
```


Status Through CLI

Run the following CLI command to get status of the distributed load balancer.

`get load-balancer <UUID_LoadBalancer> status`

```
Load Balancer
UUID : 8721fb3e-dbef-4d9a-8f48-432e893883f1
Display-Name : DLB_Service21
Status : ready
Ready LSP Count : 4
Not Ready LSP Count: 0
Partially Ready LSP Count : 0
```

Run Traceflow on Distributed Load Balancer

Run Traceflow between virtual machines or interfaces where Distributed Load Balancer is enforced and the Distributed Load Balancer virtual IP address (VIP).

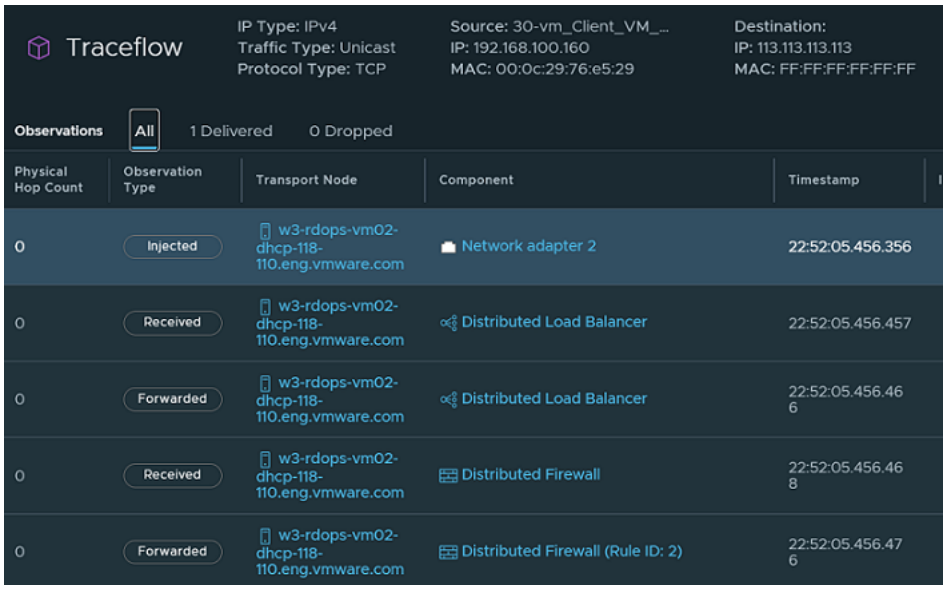
Use Traceflow for debugging purposes, when:

- On a client guest VM, where a Distributed Load Balancer service is applied, if the communication is lost between the guest VMs to Distributed Load Balancer.
- Or when guest VMs sending east-west traffic to another VM in the network drops.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Navigate to the **Plan & Troubleshoot** tab, and click **Traceflow**.
3. Enter source (for example, virtual machine where Distributed Load Balancer is enforced) and destination (Distributed Load Balancer VIP).

4 Click **Trace**.



5 View results.

- As DLB is not based on rules, if Distributed Load Balancer drops a packet, there are not rule IDs that show up in the traceflow output details.
- If Distributed Load Balancer and DFW services are applied to a client virtual machine, a member of a group, first Distributed Load Balancer service is applied followed by DFW rules service.

Supported Features

Supported features by Distributed Load Balancer (DLB).

IDPS and Anti-malware Prevention

On the **Security** → **IDS/IPS & Anti-malware** page, create a distributed firewall rule where destination is DLB pool server members and apply an IDPS security profile in the **Detect Only** or **Detect and Prevent** mode. Any traffic destined to the selected DLB pool server group is first processed by IDPS. If IDPS does not detect any malware it redirects traffic to the DLB pool servers.

vMotion

You can perform vMotion to migrate VMs from one host to another without disrupting load balancing performed by DLB. During vMotion of the VM and after migration, DLB continues to load balance traffic without dropping any packets.

Ethernet VPN (EVPN)

12

EVPN (Ethernet VPN) is a standards-based BGP control plane that provides the ability to extend Layer 2 and Layer 3 connectivity between different data centers.

This chapter includes the following topics:

- Overview of BGP EVPN
- EVPN Support in NSX
- EVPN Inline Mode
- EVPN Route Server Mode

Overview of BGP EVPN

Ethernet VPN (EVPN) is a standards-based BGP distributed control plane for Network Virtualization Overlay (NVO), that provides Layer 2 (bridging) and Layer 3 (routing) connectivity over IP or IP/MPLS underlay networks. BGP EVPN was initially designed to be used with MPLS data plane to address limitations of VPLS in service provider networks. However, EVPN has been widely adopted in data centers as a control plane mechanism for VXLAN overlay networking due to advantages in BGP scalability and flexibility.

Some of the key characteristics and benefits for BGP EVPN are:

- BGP-based control plane learning for Layer 2 and Layer 3 end host reachability information. This replaces flood-and-learn behavior of legacy L2VPN solutions, such as VXLAN, SPB, and Trill.
- ARP suppression to minimize unnecessary ARP and ND message flooding.
- Support for bridging-only and/or integrated routing and bridging.
- Support for MAC and IP mobility and multihoming.

MP-BGP EVPN Address Family

A new MP Address Family Indicator/Subsequent Address Family Indicator (AFI/SAFI) is defined for EVPN: l2vpn (25) /evpn (70). For two BGP speakers to exchange EVPN Network Layer Reachability Information (NLRI), they must negotiate the EVPN BGP capability at the start of the BGP session to ensure that both peers can support such NLRI.

Route Distinguisher and Route Targets

BGP EVPN uses the same mechanisms of other BGP VPN technologies to ensure uniqueness and multi-tenancy:

Mechanism	Description
Route Distinguisher (RD)	■ Used in EVPN to make addresses globally unique. ■ The same encoding defined in RFC 4364 is applicable for BGP EVPN. ■ RD type 0 has an Administrator subfield of 2 bytes and Assigned Number subfield of 4 bytes. The Administrator subfield must contain an Autonomous System number. ■ RD type 1 has an Administrator subfield of 4 bytes and Assigned Number subfield of 2 bytes. The Administrator subfield must contain an IP address.
Route Target (RT)	■ Used in EVPN to indicate virtual network membership by importing and exporting RTs as required.

EVPN Route Types

EVPN NLRI is further classified by the following route types:

Route Type	Description	Purpose	NSX Inline Mode	NSX Route Server Mode
Type 1 (RT-1)	Ethernet Auto-Discovery (A-D) route	Used in data centers to support EVPN active-active of multihoming.	No	Yes (receive only for DC-GW mutlihoming)
Type 2 (RT-2)	MAC/IP Advertisement route	Advertises reachability of a specific MAC address, and optionally MAC and IP address binding.	No	Yes
Type 3 (RT-3)	Inclusive Multicast Ethernet Tag route	Advertises reachability of a VNI associated to a particular VTEP in a virtual network. Type 3 routes are required for BUM traffic delivery across EVPN networks.	No	Yes
Type 4 (RT-4)	Ethernet Segment route	Used in data centers with multihomed endpoints, and used for Designated Forwarder Election to ensure that only one of the VTEPs forwards BUM traffic.	No	No
Type 5 (RT-5)	IP Prefix route	Advertises IPv4 and IPv6 prefixes reachability. This advertisement of prefixes into the EVPN domain provides the ability to build L3VPN similar services.	Yes	Yes
Type 6 (RT-6)	Selective Multicast Ethernet Tag route	Used to advertise the intent of the host or VM to receive multicast traffic for a certain Multicast Group (*,G) or Source-Group combination (S,G).	No	No

Virtual Routing and Forwarding

Virtual routing and forwarding (VRF) makes it possible to instantiate isolated routing and forwarding tables within a router. These routing and forwarding table instances isolate Layer 3 domains and segments from each other creating a multi-tenant network, either locally within the router or across multiple routers.

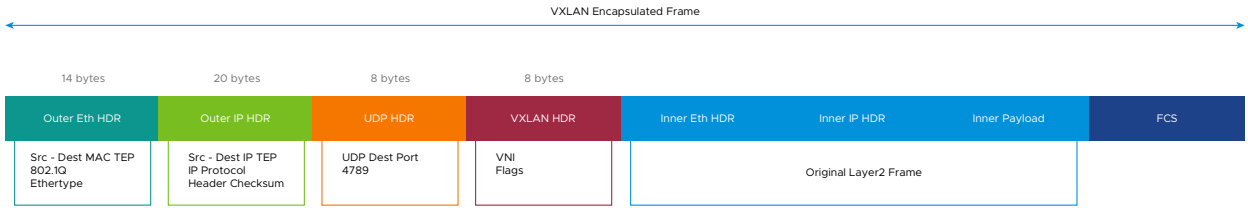
VRFs are supported in NSX by deploying tier-0 VRF gateways. A tier-0 VRF gateway must be linked to a parent tier-0 gateway and inherits some of the tier-0 settings, such as the HA mode, edge cluster, internal transit subnet, T0-T1 transit subnets, and BGP local ASN.

Multiple tier-0 VRF gateways can be created under the same parent tier-0, allowing the separation of segments and tier-1 gateways into multiple isolated tenants. With tier-0 VRF gateways, tenants can use overlapping IP addresses without any interference or communication with each other.

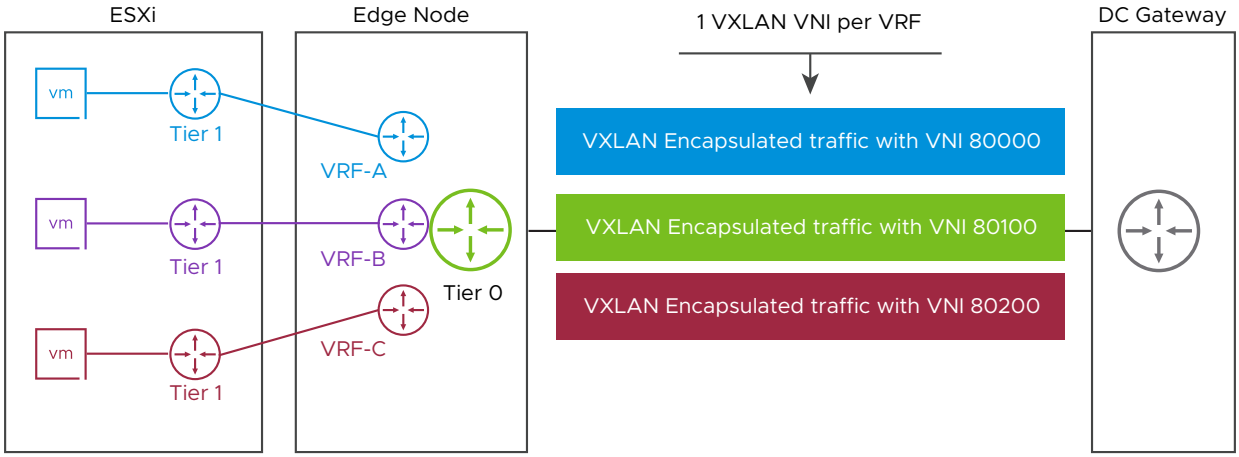
In the context of EVPN, each Layer 3 VRF is identified by a global unique Virtual Network Identifier (VNI). The VNI for each VRF must match in NSX Edge nodes and data center gateways.

VXLAN Encapsulation and VNI

VXLAN encapsulation as defined in RFC7348, is used between NSX tunnel endpoints (edge nodes for Inline mode and hypervisors for Route Server mode) and external routers in order to ensure data plane compatibility with other vendors. Inside the NSX domain, GENEVE encapsulation is still used.



The VNI is a 24-bit identifier used to identify a particular virtual network segment. When EVPN is used to advertise IP prefixes reachability by using Route Type 5 and the encapsulation type as VXLAN then the VNI identifies the tenant VRF instance. As defined in RFC9135, this VNI is advertised in the BGP control plane along with the prefix routes as well as used in the data plane encapsulation to differentiate the traffic between VRFs. The VNI for each VRF must match in NSX Edge nodes and data center gateways.



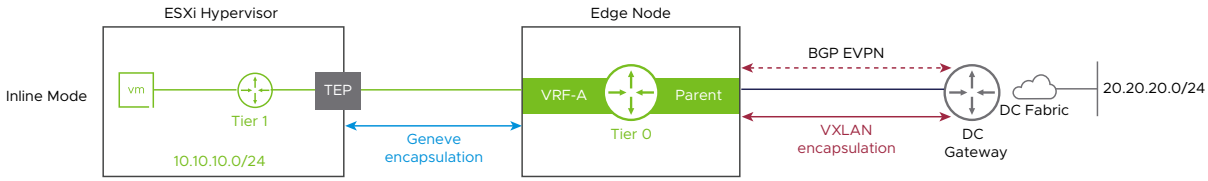
EVPN Support in NSX

NSX leverages BGP EVPN technology to interconnect and extend NSX-managed overlay networks to other data center environments not managed by NSX, VXLAN encapsulation is used between NSX TEPs (edge nodes and hypervisors) and external network devices to ensure data plane compatibility.

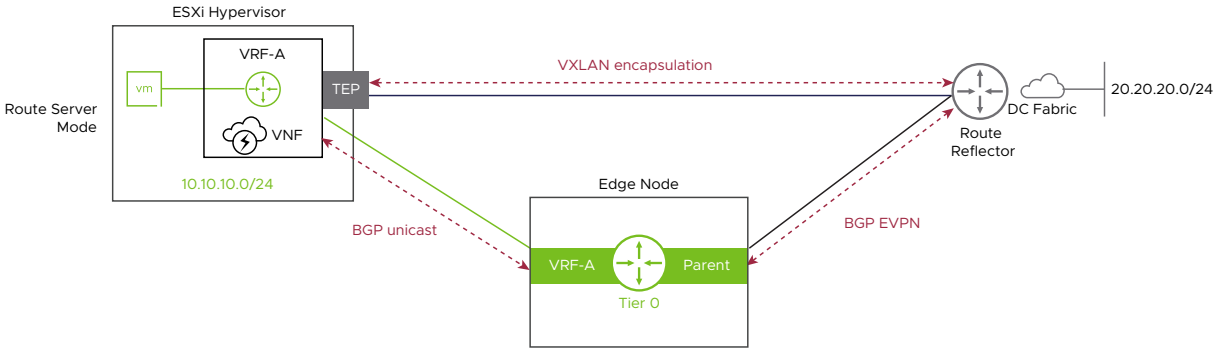
Two connectivity modes are supported for EVPN implementation in NSX:

Inline Mode

In this mode, the tier-0 gateway establishes MP-BGP EVPN control plane sessions with external routers to exchange routing information. In the data plane, edge nodes forwards all the traffic exiting the local data center to the data center gateways and incoming traffic from the remote data center to the hypervisors in the local data center. Since the edge nodes are in the data forwarding path, this model is called the Inline model.



Route Server Mode
In this mode, the tier-0 gateway establishes MP-BGP EVPN control plane to exchange routing information with the external router or route reflectors. In the data plane, ESXi hypervisors forward the traffic to external networks either to the data center gateways or remote ToR switches over VXLAN tunnels. TEPs used for the data plane VXLAN encapsulation are the same than the ones used for GENEVE encapsulation.



Route Distinguishers and Route Targets in NSX

With NSX BGP implementation, route distinguishers (RD) can be either set automatically or manually. The following table details the supported RD modes in the Inline and Route Server modes.

Mode	Auto RD	Manual RD
Inline	■ Supported. ■ Only type-1 is supported. ■ You must configure the RD Admin field. The RD Admin field must be in the format of an IP address. ■ The RD admin field is used to fill the Administrator subfield in the RD. ■ The 2-byte Assigned Number subfield will be allocated a random number in the range for each RD generation. ■ Generated auto RD is checked against other manually configured RDs to avoid any duplicates.	■ Supported. ■ Both type-0 and type-1 are allowed, but type-1 is recommended. ■ No RD Admin field is required to be configured. ■ Configure manual RD is checked against other auto RDs to avoid any duplicates.
Route Server	■ Not supported.	■ Supported. ■ Both type-0 and type-1 are allowed, but type-1 is recommended. ■ No RD Admin field is required to be configured. ■ Configured manual RD is checked against other auto RDs to avoid any duplicates.

Limitations and Caveats

- NSX supports L3 EVPN by advertising and receiving IP prefixes as EVPN Route Type-5.
- NSX generates a unique route MAC for every NSX Edge VTEP in the EVPN domain. However, there may be other nodes in the network that are not managed by NSX, for example, physical routers. You must make sure that the router MACs are unique across all the VTEPs in the EVPN domain.
- The EVPN feature supports NSX Edge nodes to be either the ingress or the egress of the EVPN virtual tunnel endpoint. If an NSX Edge node receives EVPN Route Type-5 prefixes from its eBGP peer that needs to be redistributed to another eBGP peer, the routes are re-advertised without any change to the next hop.
- In multi-path network topologies, it is recommended that ECMP is enabled for the NSX BGP EVPN control plane, so that all the possible paths can be advertised by the tier-0 gateway. This will avoid any potential traffic blackhole due to asymmetric data path forwarding.
- A tier-0 gateway can span across multiple edge nodes. However, specifying a unique route distinguisher for each edge node or TEP (either via auto or manual configuration) is not supported. As a result, the use of ECMP on the peer router is not supported.
- Route maps are not supported for EVPN address family.
- Recursive route resolution for gateway IP via default static route is not supported.

Limitations and caveats for Inline mode:

- Only BGP Graceful Restart in Helper Mode is supported.
- Only eBGP is supported between tier-0 SRs and external routers.
- Only one TEP is supported per edge node. The use of loopback interfaces for TEP is highly recommended.

Limitations and caveats for Route Server mode:

- The High Availability mode on the tier-0 must be set to active-active.
- The manual Route Distinguisher and manual Route Targets are supported.
- BGP Graceful Restart, Helper Mode, and Restarted Mode are not supported.
- Only eBGP is supported between hosted VNFs and tier-0 VRF gateways.
- eBGP multihop using loopbacks is required between tier-0 SRs and external routers. Using uplinks for eBGP neighbor session is not supported for EVPN Router Server mode operation.
- The VNF uplink towards the tier-0 SR VRF must be in the same subnet as the Integrated Routing and Bridging (IRB) on the data center gateways.

EVPN Inline Mode

In the EVPN Inline mode, MP-BGP sessions with the L2VPN EVPN address family are configured between tier-0 gateways and the external routers. The tier-0 gateway will negotiate the L2VPN EVPN AFI/SAFI with the external BGP peer (data center gateway) and start exchanging EVPN information. In this mode, the edge nodes are in the datapath between internal workloads and external networks.

The NSX EVPN Inline mode is based on the "Interface-less IP-VRF-to-IP-VRF Model" as defined in the IETF RFC9136. In this mode, the tier-0 gateway advertises to the data center gateway only Route Type 5 (RT-5) routes, that includes:

- Prefixes of segments connected directly to the tier-0 VRF gateway.
- Tier-1 segment prefixes redistributed to the tier-0 VRF gateway.
- Other redistributed sources, such as static routes and connected interfaces.

- BGP prefixes learned via southbound BGP sessions inside the VRF.

All routes received from and advertised to the data center gateways as RT-5 has an EVPN Router's MAC Extended Community with the MAC address of the corresponding peer uplink. Since "Interface-less IP-VRF-to-IP-VRF Model" is used, there is no recursive route lookup to resolve the RT-5 route. The packet is encapsulated in VXLAN using the RT-5's next hop as destination IP address and EVPN Router's MAC Extended Community as MAC address.

Inside the NSX domain, the control plane is still handled by the central control plane (CCP) and the encapsulation protocol among internal TEPs is still GENEVE.

EVPN Inline Mode Configuration Workflow

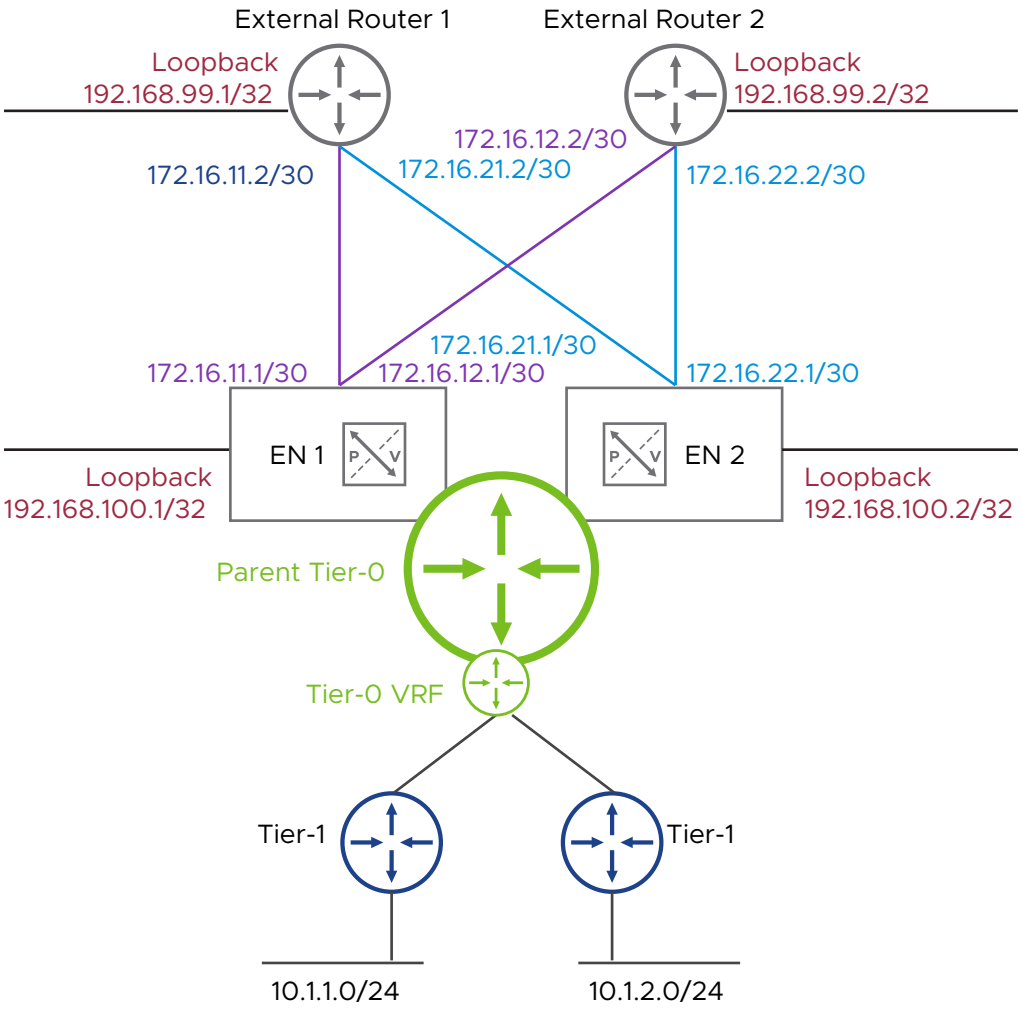
Follow this workflow to configure EVPN with Inline mode.

Prerequisites

A typical BGP EVPN Inline mode deployment topology has the following characteristics:

- There are point-to-point uplinks between edge nodes and external routers over individual VLAN segments.
- There are BGP peering sessions between edge nodes and external routers using loopback interfaces.
- Loopback reachability can be achieved using either static routing or OSPF protocol.

The following diagram depicts a typical BGP EVPN Inline mode deployment topology:



Procedure

1 Create a VNI pool. See Add an EVPN/VXLAN VNI Pool.

2 Configure a tier-0 gateway and enable EVPN. See Configure a Tier-0 Gateway for EVPN Inline Mode.

3 Configure BGP neighbors. See Configure BGP Neighbors for a Tier-0 Gateway.

4 Configure a tier-0 VRF gateway. See Configure a Tier-0 VRF Gateway for EVPN Inline Mode.

5 Verify the BGP neighbor session status.

   a Select **Networking > Tier-0 Gateways**.

   b Click the menu icon (three dots) of the tier-0 gateway and select **Generate BGP Summary**.

   c Verify that **Connection Status** for the neighbor is **Established** and that **Address Families** displays **L2VPN EVPN**.

6 Verify the tier-0 VRF gateway forwarding table.

   a Select **Networking > Tier-0 Gateways**.

   b Click the menu icon (three dots) of the tier-0 VRF gateway and select **Download Forwarding Table**.

   c Verify that the remote routes received from the external router are installed in the tier-0 VRF gateway forwarding table.

Configure a Tier-0 Gateway for EVPN Inline Mode

Use this procedure to configure a tier-0 gateway for EVPN Inline mode.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Networking > Tier-0 Gateways**.

3 Create a tier-0 gateway and configure uplink interfaces between the tier-0 gateway and data center gateways. See Add a Tier-0 Gateway.

4 (Optional) Configure loopback interfaces for each edge node to be used as the BGP source address and VXLAN VTEPs.

5 (Optional) If connecting to BGP peers using loopbacks, make sure that the data center gateway external loopbacks are reachable from the parent tier-0 gateway.

   This connectivity can be achieved either by using OSPF or static routes.

6 Enable EVPN.

   a Expand the **EVPN Settings** section.

   b For **EVPN Mode**, select **Inline**.

   c For **EVPN/VXLAN VNI Pool**, select a VNI pool.

   d For **EVPN Tunnel Endpoint**, click **Set > Add EVPN Local Tunnel Endpoint**.

      1 Enter a name for the endpoint.

      2 Select the edge node.

      3 Enter the IP address for the VXLAN TEP.

      4 Click **Save** and then **Close**.

7 (Optional) If the VXLAN TEP IP address is different than the loopback interface IP address used for the BGP neighbor configuration, make sure that the local tunnel endpoints is redistributed into the tier-0 BGP routing table.

   a Expand the **Route Re-distribution** section.

   b Click **Set > Add Route Re-distribution**.

      1 Enter a name for the re-distribution policy.

      2 For **Destination Protocol**, select **BGP**.

      3 Click **Set** and then select **EVPN TEP IP**.

      4 Click **Apply**.

      5 Click **Add** and then **Apply**.

Configure BGP Neighbors for a Tier-0 Gateway

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Networking > Tier-0 Gateways**.

3 Click the menu icon (three dots) for the tier-0 gateway and select **Edit**.

4 Expand the **BGP** section.

5 Click the **BGP** toggle to enable BGP.

6 Enter the **Local AS** number.

7 Configure BGP neighbors.

   a For **BGP Neighbors**, click **Set > Add BGP Neighbor**.

   b Enter the IP address of the neighbor.

   c Enable BFD if required.

   d Enter the **Remote AS number** of the neighbor.

   e For **Source Addresses**, enter the source IP address.

      There should be one or more addresses of created external interfaces or loopback.

   f For **Route Filter**, click **Set > Add Route Filter**.

      1 For **IP Address Family**, select **L2VPN EVPN**.

      2 Specify the desired maximum routes.

      3 Click **Add** and then **Apply**.

8 Click **Save** and then **Close**.

Configure a Tier-0 VRF Gateway for EVPN Inline Mode

Use this procedure to configure a tier-0 VRF gateway for EVPN Inline mode.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Networking > Tier-0 Gateways**.

3 Add a VRF gateway.

   a Click **Add Gateway > VRF**.

   b Enter a name for the gateway.

   c Select a tier-0 gateway with EVPN Inline mode enabled to connect to.

   Some advanced configurations are inherited from the parent tier-0 gateway, such as HA mode, edge cluster, internal transit subnet, T0-T1 transit subnets, and BGP Local ASN.

4 Expand the **VRF Settings** section.

   a Enter the route distinguisher.

   If the connected parent tier-0 gateway has **RD Admin Address** configured, then the **Route Distinguisher** field is automatically populated. Enter a new value if you want to override the assigned route distinguisher.

   b For **Route Targets**, click **Set > Add Route Target** to add route targets.

   - For EVPN, only the **Manual** mode is supported.
   - Specify one or more **Import Route Targets** to install the desired received routes from BGP peers to the VRF routing table.
   - Specify one or more **Export Route Targets** to label advertised VRF routes.

   c Click **Add** and then **Apply**.

   d For **EVPN Transit VNI**, enter the L3 transit VNI value for the VRF.

   This VNI must be unique per VRF and belong to the configured EVPN/VXLAN VNI pool. Make sure that the same VNI value for the VRF is configured on the external router.

5 Click **Save** and then **Yes** to continue configuring the VRF gateway.

6 Expand the **Route Re-distribution** section.

   a Click **Set > Add Route Re-Distribution**.

   b Enter a name for the re-distribution policy.

   c Click **Set** to select available sources, such as tier-0 connected interfaces and segments.

   d Click **Apply**.

   e Click **Add** and then **Apply**.

7 Make sure that the created segments or tier-1 gateways are connected to the new tier-0 VRF gateway.

EVPN Route Server Mode

In the EVPN Route Server mode, the tier-0 service router (SR) hosted on the edge node acts as a BGP route server, establishing BGP control plane sessions with southbound VNFs and external data center routers. ESXi hypervisors exchange the user plane traffic directly with the data center fabric routers using VXLAN encapsulation, bypassing the edge node in the data path.

From the BGP control plane perspective, there are two types of sessions:

Session Type	Description
Between hosted VNFs and tier-0 VRF gateway.	■ BGP IPv4 unicast and IPv6 unicast sessions from the VNF to the tier-0 VRF service ports. ■ IP prefixes learned from the VNF via the BGP IPv4/IPv6 unicast sessions are advertised as EVPN Route-Type 5 towards the external route with the corresponding VRF route distinguisher and route targets. ■ IP prefixes (RT-5) learned from the external router via the BGP EVPN session are injected in the tier-0 VRF routing table based on the route target policies and are advertised as IPv4/IPv6 unicast routes to the VNF.
Between tier-0 SR and DC gateways.	■ BGP IPv4 session with L2VPN EVPN address family from tier-0 SR to the loopback of DC gateways. ■ The IP prefixes learned from the VNF via the BGP IPv4/IPv6 unicast sessions are advertised as EVPN Route-Type 5 towards the external router with the corresponding VRF route distinguisher and route targets. ■ IP prefixes (RT-5) learned from the external router via the BGP EVPN session are injected in the tier-0 VRF routing table based on the route target policies and are advertised as IPv4/IPv6 unicast routes to the VNF. ■ EVPN routes of type 2 (RT-2) are exchanged between tier-0 SR and DC gateways to advertise MAC and IP/MAC bindings for RT-5 routes next hops. ■ EVPN routes of type 1 (RT-1) are sent by the DC gateway to the tier-0 SRs when multiple data center gateways are configured as L2ECMP EVPN multihoming. RT-1 routes are used to create a list of data center gateway VTEPs to be used as next hop by ESXi nodes.

The NSX EVPN Route Server mode is based on the "Interface-ful IP-VRF-to-IP-VRF with SBD IRB" as defined in the IETF RFC 9136. The RFC 9136 introduces a concept called overlay index in EVPN. A key concept of EVPN RT-5 is the overlay index, which can be a gateway IP address, a MAC, or an ESI. When a node receives an EVPN RT-5 with an overlay index specified, the receiving node performs a recursive route resolution to find the appropriate node to forward the data packets for the corresponding IP prefix.

NSX EVPN Route Server mode implements the gateway IP address as the overlay index. The tier-0 SR also advertises to the external router an additional EVPN type-2 route with the appropriate MAC/IP (gateway IP) binding and the corresponding VXLAN TEP address.

The gateway IP address in this case will be the IPv4 BGP next hop for a given prefix as advertised by the VNF to the tier-0 VRF gateway. The recursive route resolution uses respective RT-2 to learn the ESXi TEP address where the VNF is hosted.

Data Center Gateway Requirements

The data center gateway router connected to the edge node tier-0 SR must support the "Interface-ful IP-VRF-to-IP-VRF with SBD IRB Mode" described in the IETF RFC 9136, section 4.4.2.

Virtual Network Function (VNF) Requirements

A VNF is typically a virtual machine used for some networking function such as a virtual router, firewall, or a Telco 5G core application. In the context of EVPN Route Server mode, the VNF is hosted by an ESXi hypervisor and should support 802.1Q-tagged interfaces and regular BGP protocol with IPv4 and IPv6 unicast address families.

EVPN Route Server Mode Configuration Workflow

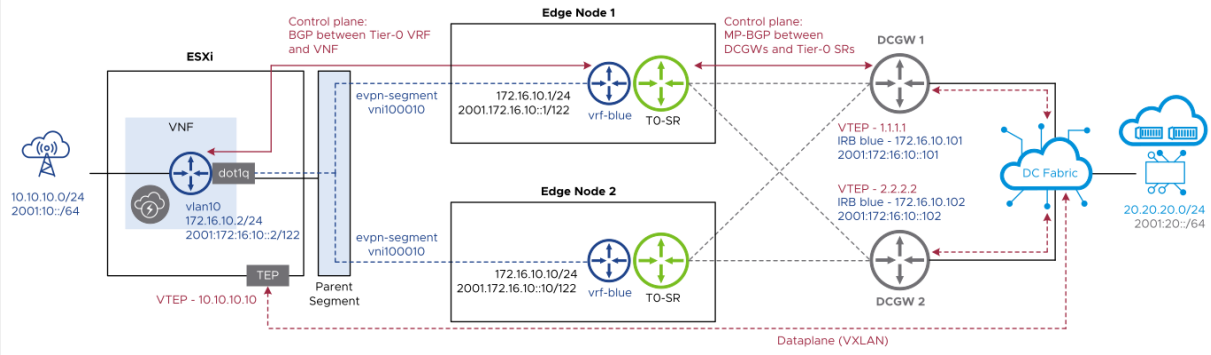
Follow this workflow to configure EVPN with Route Server mode.

Prerequisites

A typical BGP EVPN Route Server mode deployment topology has the following characteristics:

- The tier-0 gateway must be in active-active mode.
- There are least two data center gateways connected to the edge nodes.
- There are point-to-point uplinks between edge nodes and data center gateways over VLAN segments.
- There are eBGP peering sessions between edge nodes and data center gateways using loopback interfaces.
- The ESXi node TEP network must have connectivity to the data center gateway VTEP IP addresses.
- There are southbound VMs and workloads connected to the VNF southbound interfaces using regular NSX segments.
- There are eBGP peering sessions between the VNF and service ports of tier-0 VRF gateways.

The following diagram depicts a typical BGP EVPN Route Server mode deployment topology:



Procedure

1. Configure a VNI pool. See Add an EVPN/VXLAN VNI Pool.
2. Configure an EVPN tenant. See Configure an EVPN Tenant.
3. Configure EVPN BFD between the ESXi nodes and data center gateways. See Configure EVPN BFD.
4. Configure a tier-0 gateway and enable EVPN. See Configure a Tier-0 Gateway for EVPN Route Server Mode.
5. Configure BGP neighbors. See Configure BGP Neighbors for a Tier-0 Gateway.
6. Configure a tier-0 VRF gateway. See Configure a Tier-0 VRF Gateway for EVPN Route Server Mode.
7. Configure networking for onboarding the tenant VNF. See Onboard a Tenant VNF for EVPN Route Server Mode.

8 In NSX, verify the following:

Verification	Steps
Verify the tier-0 SR BGP neighbor session status.	1 Select **Networking > Tier-0 Gateways**. 2 Click the menu icon (three dots) for the tier-0 VRF gateway and select **Generate BGP Summary**. 3 Verify that **Connection Status** displays **Established**. 4 Verify that **Address Families** displays **L2VPN EVPN**.
Verify the tier-0 VRF BGP neighbor session status.	1 Select **Networking > Tier-0 Gateways**. 2 Click the menu icon (three dots) for the tier-0 VRF gateway and select **Generate BGP Summary**. 3 Verify that the neighbor **Connection Status** displays **Established**.
Verify the tier-0 VRF gateway routing table.	1 Select **Networking > Tier-0 Gateways**. 2 Click the menu icon (three dots) for the tier-0 VRF gateway and select **Download Routing Table**. 3 Select the transport node (edge node) and for **Source**, select **BGP**. 4 Click **Download**. 5 Verify that the remote nodes received from the external router are installed in the tier-0 VRF gateway routing table.

9 In the ESXi NSX CLI, verify the following.

- Verify the status of the VTEP group for the VNI.

```
get logical-switch <vni> vtep-group
```

- Verify the MAC address for the L2 VNI.

```
get logical-switch <vni> mac-table
```

- Verify the ARP table for the L2 VNI.

```
get logical-switch <vni> arp-table
```

Configure an EVPN Tenant

Configuring an EVPN tenant is required for EVPN Route Server mode.

To configure an EVPN tenant, you must specify one or more VLAN-VNI mappings.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Click **Networking > EVPN Tenant > Add EVPN Tenant**.

3 Enter a name.

4 For **Overlay Transport Zone**, select one from the drop-down list.

5 Select an existing EVPN/VXLAN VNI pool or create a new pool by clicking the menu icon (three dots).

6 For **VLAN-VNI Mapping**, click **Set** to create one or more mappings.

The VNIs must be within the range specified in the EVPN/VXLAN VNI pool. You can specify a single value for VLAN and VNI, or a range of values such as 10-20 for VLAN and 77010-77020 for VNI. For each VLAN-VNI mapping specified in the EVPN tenant, a special segment for the VNI will be created automatically. You can see the VRF segments by navigating to **Networking > Segments**.

7 Click **Save**.

Configure EVPN BFD

In EVPN Server Route mode, BFD (Bidirectional Flow Detection) between the ESXi nodes and data center gateways can be "optionally" enabled for TEP fast failure detection.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Networking > Global Networking Config**.

3 Click **Edit**.

4 For **EVPN BFD Profile**, select a BFD profile.

If you need to create a new profile with required BFD timers, click the menu icon (three dots) and select **Create New**.

5 For **Enable EVPN BFD**, make sure that the toggle is on.

6 Click **Save**.

Configure a Tier-0 Gateway for EVPN Route Server Mode

Use this procedure to configure a tier-0 gateway for EVPN Route Server mode.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Networking > Tier-0 Gateways**.

3 Create a tier-0 gateway and configure uplink interfaces between the tier-0 gateway and data center gateways. See Add a Tier-0 Gateway.

Note For EVPN Route Server mode, only active-active HA mode is supported.

4 (Optional) Configure loopback interfaces for each edge node to be used as the BGP source address and VXLAN VTEPs.

5 Make sure that the data center gateway external loopbacks are reachable from the parent tier-0 gateway.

This connectivity can be achieved either by using OSPF or static routes.

6 Enable EVPN.
   a Expand the **EVPN Settings** section.
   b For **EVPN Mode**, select **Route Server**.
   c Select the EVPN tenant.
   d Click **Save** and then **Close**.

Configure BGP Neighbors for a Tier-0 Gateway

Procedure

1 With admin privileges, log in to NSX Manager.
2 Select **Networking > Tier-0 Gateways**.
3 Click the menu icon (three dots) for the tier-0 gateway and select **Edit**.
4 Expand the **BGP** section.
5 Click the **BGP** toggle to enable BGP.
6 Enter the **Local AS** number.
7 Configure BGP neighbors.
   a For **BGP Neighbors**, click **Set > Add BGP Neighbor**.
   b Enter the IP address of the neighbor.
   c Enable BFD if required.
   d Enter the **Remote AS number** of the neighbor.
   e For **Source Addresses**, enter the source IP address.

      There should be one or more addresses of created external interfaces or loopback.
   f For **Route Filter**, click **Set > Add Route Filter**.
      1 For **IP Address Family**, select **L2VPN EVPN**.
      2 Specify the desired maximum routes.
      3 Click **Add** and then **Apply**.
8 Click **Save** and then **Close**.

Configure a Tier-0 VRF Gateway for EVPN Route Server Mode

Use this procedure to configure a tier-0 VRF gateway for EVPN Route Server mode.

Procedure

1 With admin privileges, log in to NSX Manager.
2 Select **Networking > Tier-0 Gateways**.

3 Add a VRF gateway.

a Click **Add Gateway > VRF**.

b Enter a name for the gateway.

c Select a tier-0 gateway with EVPN Route Server mode enabled to connect to.

Some advanced configurations are inherited from the parent tier-0 gateway, such as HA mode, edge cluster, internal transit subnet, T0-T1 transit subnets, and BGP Local ASN.

4 Expand the **VRF Settings** section and configure the L3 VRF settings.

Note The L3 VRF settings correspond with the EVPN RT-5 routes advertised and received from the data center gateways. The imported/exported RTs needs to match with the corresponding L3 VRF configuration in the data center gateway side.

a Enter the route distinguisher.

If the connected tier-0 gateway has **RD Admin Address** configured, the **Route Distinguisher** field is automatically populated. Enter a new value if you want to override the assigned route distinguisher.

b For **Route Targets**, click **Set > Add Route Target** to add route targets.

- For EVPN, only the **Manual** mode is supported.
- Specify one or more **Import Route Targets** to install the desired routes from BGP peers to the VRF routing table.
- Specify one or more **Export Route Targets** to label advertised VRF routes.

c Click **Add** and then **Apply**.

5 Configure the L2 VNI settings.

Note The L2 VNI settings correspond with the EVPN RT-3, RT-2, and RT-1 routes advertised and received from the data center gateways. The imported/exported RTs needs to match with the corresponding L2 VNI configuration in the data center gateway side.

a For **L2 VNI**, click **Set > Add L2 VNI**.

b Select an L2 VNI.

The VNI must be unique and belong to the VLAN-VNI mappings from an EVPN tenant.

c Enter the route distinguisher.

d For **Route Targets**, click **Set > Add**.

- For EVPN, only the **Manual** mode is supported.
- Specify one or more **Import Route Targets** to install the desired routes from BGP peers to the VRF routing table.
- Specify one or more **Export Route Targets** to label advertised VRF routes.

e Click **Add** and then **Apply**.

f For **VTEP Groups**, click the toggle to enable or disable the creation of VTEP groups with data center gateways for the L2 VNI.

Note The **VTEP Groups** option will enable northbound L2ECMP EVPN multihoming between the VNF and DC gateways.

6 Click **Save** and then **Yes** to continue configuring the VRF gateway.

7 Expand the **Route Re-distribution** section.

a Click **Set > Add Route Re-Distribution**.

b Enter a name for the re-distribution policy.

c Click **Set** to select available sources, such as tier-0 connected interfaces and segments.

d Click **Apply**.

e Click **Add** and then **Apply**.

Onboard a Tenant VNF for EVPN Route Server Mode

For EVPN Route Server mode, you need to onboard a virtual network function (VNF). A VNF is typically a virtual machine used for some networking function such as a virtual router, firewall, or a Telco 5G core application. In the context of EVPN Route Server mode, the VNF is hosted by an ESXi hypervisor and should support 802.1Q-tagged interfaces and regular BGP protocol with IPv4 and IPv6 unicast address families.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Create a parent overlay segment to connect the VNF virtual machine.

a Select **Networking > Segments**.

b Click **Add Segment**.

c Enter a name for the segment.

d For **Connected Gateway**, select **None**.

e Select an overlay transport zone.

f For the EVPN configuration section, select the EVPN tenant.

g Click **Save**.

3 Create a service interface to establish a BGP IPv4 session between the tier-0 VRF gateway and the VNF.

a Expand the **Interfaces** section.

b Click **Set > Add Interface**.

c Enter a name for the interface.

d For **Type**, select **External**.

e Enter an IP address.

f Select a segment to connect to.

The segment should be one of the EVPN automatically created segments with the appropriate VLAN to communicate with the VNF.

g Select an edge node.

h Click **Save** and then **Close**.

4 In VMware vCenter, connect the VNF virtual machine uplink interface to the NSX parent segment created from the previous step.

5 Link the VNF segment port to the corresponding EVPN VLAN.

a In NSX, select **Networking > Segments**.

b Click the menu icon (three dots) for the parent segment and select **Edit**.

c For **Ports / Interfaces**, click **Set**.

For each VNF interface attached to the parent segment, you should see the corresponding segment port.

d Click the menu icon (three dots) of the segment port and select **Edit**.

e For **EVPN VLAN**, add the corresponding VLAN.

The VLAN should match the VLAN/VNI mapping for the VRF.

f Click **Save** and then **Close**.

6 Configure the BGP IPv4 session between the tier-0 VRF gateway and the VNF.

a Select **Networking > Tier-0 Gateways**.

b Click the menu icon (three dots) of the tier-0 VRF gateway and select **Edit**.

c Expand the **BGP** section.

d For **BGP**, click the toggle to enable BGP.

e You can configure advanced BGP settings, such as ECMP.

f For **BGP Neighbors**, click **Set > Add BGP Neighbor**.

1 Enter the neighbor IP address.

2 For **BFD**, click the toggle to enable or disable the BFD session with the VNF.

3 Enter the remote AS number of the neighbor.

4 For **Source IP Address**, it is not required. The system automatically uses the service port interface IP address previously created.

5 For **Route Filter**, click **Set > Add Route Filter** to enable IP address families and the desired maximum routes.

6 For **IP Address Family**, select **IPv4** or **IPv6**.

7 Click **Add** and then **Apply**.

g Expand the **Timers & Password** section.

h Configure the BFD timers and BGP password.

i Click **Save** and then **Close**.

Forwarding Policies

13

This feature pertains to NSX Cloud.

Forwarding Policies or Policy-Based Routing (PBR) rules define how NSX handles traffic from an NSX-managed VM. This traffic can be steered to NSX overlay or it can be routed through the cloud provider's (underlay) network.

Note See Chapter 28 Using NSX Cloud for details on how to manage your public cloud workload VMs with NSX.

Three default forwarding policies are set up automatically after you either deploy a PCG on a Transit VPC/VNet or link a Compute VPC/VNet to the Transit.

- **Route to Underlay** for all traffic that is addressed within the Transit/Compute VPC/VNet
- **Route from Underlay** for all traffic destined to the metadata services of the public cloud.
- **Route to Overlay** for all other traffic, for example, traffic that is headed outside the Transit/Compute VPC/VNet. Such traffic is routed over the NSX overlay tunnel to the PCG and further to its destination.

  Note **For traffic destined to another VPC/VNET managed by the same PCG:** Traffic is routed from the source NSX-managed VPC/VNet via the NSX overlay tunnel to the PCG and then routed to the destination VPC/VNet.

  For traffic destined to another VPC/VNet managed by a different PCG: Traffic is routed from one NSX-managed VPC/VNet over the NSX overlay tunnel to the PCG of the source VPC/VNet and forwarded to the PCG of the destination NSX-managed VPC/VNet.

  If traffic is headed to the internet, the PCG routes it to the destination in the internet.

Micro-segmentation while Routing to Underlay

Micro-segmentation is enforced even for workload VMs whose traffic is routed to the underlay network.

If you have direct connectivity from an NSX-managed workload VM to a destination outside the managed VPC/VNet and want to bypass the PCG, set up a forwarding policy to route traffic from this VM via underlay.

When traffic is routed through the underlay network, the PCG is bypassed and therefore the north-south firewall is not encountered by traffic. However, you still have to manage rules for east-west or distributed firewall (DFW) because those rules are applied at the VM-level before reaching the PCG.

Supported Forwarding Policies and Common Use Cases

You may see a list of forwarding policies in the drop-down menu but in this release only the following forwarding policies are supported:

- Route to Underlay
- Route from Underlay
- Route to Overlay

These are the common scenarios where forwarding policies are useful:

- **Route to Underlay**: Access a service on underlay from an NSX-managed VM. For example, access to the AWS S3 service on the AWS underlay network.
- **Route from Underlay**: Access a service hosted on an NSX-managed VM from the underlay network. For example, access from AWS ELB to the NSX-managed VM.

This chapter includes the following topics:

- Add or Edit Forwarding Policies

Add or Edit Forwarding Policies

You can edit the auto-created forwarding policies or add new ones.

Starting in NSX 3.1.1, you can utilize AWS Transit Gateway for micro-segmentation of your workload VMs in your VPCs after you create the following forwarding policy to route all your networking traffic using the AWS underlay network.

Table 13-1. Set up this Forwarding Policy if you are using AWS Transit Gateway

Option	Value
Sources	A Group or a set of groups in NSX Manager containing all NSX-managed VMs from your Transit and Compute VPCs connected using the AWS Transit Gateway.
Destinations	All (0.0.0.0/0)
Services	Any
Action	**Route to Underlay**

See more details at *Using PCG with AWS Transit Gateway* in the *NSX Installation Guide*.

The following settings are explained through the example use case: forwarding policy for services provided by the public cloud, such as S3 by AWS. Create a policy to allow a set of IP addresses to access this service by being routed through underlay.

Prerequisites

You must have a VPC or VNet with a PCG deployed on it.

Procedure

1. Click **Add Section**. Name the section appropriately, for example, **`AWS Services`**.
2. Select the check box next to the section and click **Add Rule**. Name the rule, for example, **`S3 Rules`**.
3. In the **Sources** tab, select the VPC or VNet where you have the workload VMs to which you want to provide the service access, for example, the AWS VPC. You can also create a **Group** here to include multiple VMs matching one or more criteria.
4. In the **Destinations** tab, select the VPC or VNet where the service is hosted, for example, a **Group** that contains the IP address of the S3 service in AWS.
5. In the **Services** tab, select the service from the drop-down menu. If the service does not exist, you can add it. You can also leave the selection to **Any** because you can provide the routing details under **Destinations**.
6. In the **Action** tab, select how you want the routing to work, for example, select **Route to Underlay** if setting up this policy for the AWS S3 service.
7. Click **Publish** to finish setting up the Forwarding Policy.

IP Address Management (IPAM)

14

To manage IP addresses, you can configure DNS (Domain Name System), DHCP (Dynamic Host Configuration Protocol), IP address pools, and IP address blocks.

Note IP blocks are used by NSX Container Plugin (NCP). For more info about NCP, see the *NSX Container Plug-in for Kubernetes and Cloud Foundry - Installation and Administration Guide*.

This chapter includes the following topics:

- Add a DNS Zone
- Add a DNS Forwarder Service
- Add an IP Address Pool
- Add an IP Address Block

Add a DNS Zone

You can configure DNS zones for your DNS service. A DNS zone is a distinct portion of the domain name space in DNS.

When you configure a DNS zone, you can specify a source IP for a DNS forwarder to use when forwarding DNS queries to an upstream DNS server. If you do not specify a source IP, the DNS query packet's source IP will be the DNS forwarder's listener IP. Specifying a source IP is needed if the listener IP is an internal address that is not reachable from the external upstream DNS server. To ensure that the DNS response packets are routed back to the forwarder, a dedicated source IP is needed. Alternatively, you can configure SNAT on the logical router to translate the listener IP to a public IP. In this case, you do not need to specify a source IP.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Select **Networking > IP Management > DNS**.
3. Click the **DNS Zones** tab.

4 To add a default zone, select **Add DNS Zone > Add Default Zone**

   a Enter a name and optionally a description.

   b Enter the IP address of up to three DNS servers.

   c (Optional) Enter an IP address in the **Source IP** field.

5 To add an FQDN zone, select **Add DNS Zone > Add FQDN Zone**

   a Enter a name and optionally a description.

   b Enter a FQDN for the domain.

   c Enter the IP address of up to three DNS servers.

   d (Optional) Enter an IP address in the **Source IP** field.

6 Click **Save**.

Add a DNS Forwarder Service

You can configure a DNS forwarder to forward DNS queries to external DNS servers.

Before you configure a DNS forwarder, you must configure a default DNS zone. Optionally, you can configure one or more FQDN DNS zones. Each DNS zone is associated with up to 3 DNS servers. When you configure a FQDN DNS zone, you specify one or more domain names. A DNS forwarder is associated with a default DNS zone and up to 5 FQDN DNS zones. When a DNS query is received, the DNS forwarder compares the domain name in the query with the domain names in the FQDN DNS zones. If a match is found, the query is forwarded to the DNS servers specified in the FQDN DNS zone. If a match is not found, the query is forwarded to the DNS servers specified in the default DNS zone.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Networking > IP Management > DNS**.

3 Click **Add DNS Service**.

4 Enter a name.

5 Select a tier-0 or tier-1 gateway.

6 Enter the IP address of the DNS service.

   Clients send DNS queries to this IP address, which is also known as the DNS forwarder's listener IP.

7 Select a default DNS zone.

8 Select up to five FQDN zones.

9 Select a log level.

10 Enter a description.

11 Click the **Admin Status** toggle to enable or disable the DNS service.

12 (Optional) Change the cache size.

13 Click **Save**.

Add an IP Address Pool

You can configure IP address pools for use by components such as DHCP.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Networking > IP Management > IP Address Pools**.

3 Click **Add IP Address Pool**.

4 Enter a name and optionally a description.

5 Click **Set** in the **Subnets** column to add subnets.

6 To specify an address block, select **Add Subnet > IP Block**.

   a Select an IP block.

   b Specify a size.

   c Specify the allocation range.

   Note

   - For IPv6, you must specify the allocation range.
   - The allocation range value cannot be more than 1048576.

   d Click the **Auto Assign Gateway** toggle to enable or disable automatic gateway IP assignment.

   e Click **Add**.

7 To specify IP ranges, select **Add Subnet > IP Ranges**.

   a Enter IPv4 or IPv6 IP ranges.

   b Enter IP ranges in CIDR format.

   c Enter an address for **Gateway IP**.

   d Click **Add**.

8 Click **Save**.

Note After you add an IP address pool and IP addresses have been allocated from the pool, you will not be able to delete the pool. If you want to expand the pool, you must add new IP ranges. For example, if the existing range is 192.168.1.11 - 192.168.1.20 and you want to expand it to 192.168.1.10 - 192.168.1.30, add the following two ranges:

- 192.168.1.10 - 192.168.1.10
- 192.168.1.21 - 192.168.1.30

Add an IP Address Block

You can configure IP address blocks for use by other components.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Networking > IP Management > IP Address Pools**.

3 Click the **IP Address Blocks** tab.

4 Click **Add IP Address Block**.

5 Enter a name and optionally a description.

6 Enter an IP block in CIDR format.

7 Click **Save**.

Networking Settings

15

You can configure networking settings for IPv6, VNI (Virtual Network Identifier) pools, gateways, multicast, and BFD (Bidirectional Forwarding Detection).

This chapter includes the following topics:

- Configuring Multicast
- Add an EVPN/VXLAN VNI Pool
- Configure Global Gateway Settings
- Add a Gateway QoS Profile
- Add a BFD Profile
- Add a DHCP Profile

Configuring Multicast

You can configure multicast on a tier-0 gateway and optionally on a tier-1 gateway for an IPv4 network to send the same multicast data to a group of recipients. In a multicast environment, any host, regardless of whether it is a member of a group, can send to a group. However, only the members of a group will receive packets sent to that group.

The multicast feature has the following capabilities and limitations:

- PIM Sparse Mode with IGMPv2.
- No Rendezvous Point (RP) or Bootstrap Router (BSR) functionality on NSX. However, RP information can be learned via PIM Bootstrap Messages (BSMs). In addition, multiple Static RPs can be configured.

  When a Static RP is configured, it serves as the RP for all multicast groups (224/4). If candidate RPs learned from BSMs advertise candidacy for the same group range, the Static RP is preferred. However, if candidate RPs advertise candidacy for a specific group or range of groups, they are preferred as the RP for those groups.
- The Reverse Path Forwarding (RPF) check for all multicast-specific IPs (senders of data traffic, BSRs, RPs) requires that a route to each of them exists.

- The RPF check requires a route to each multicast-specific IP with an IP address as the next hop. Reachability via device routes, where the next hop is an interface index, is not supported.
- Both tier-0 and tier-1 gateways are supported. To enable multicast on a tier-1 gateway, an Edge cluster must be selected and the tier-1 gateway must be linked to a tier-0 gateway that also has multicast enabled.
- All uplinks on a tier-0 gateway are supported.
- Multiple Static RPs with discontinuous group ranges are supported.
- IGMP local groups on uplink interfaces are supported.
- PIM Hello Interval and Hold Time are supported.
- (NSX 3.2.0) Active-Cold Standby only is supported. The NSX Edge cluster can be in active-active or active-standby mode. When the cluster is in active-active mode, two of the cluster members will run multicast in active-cold standby mode. You can run the CLI command `get mcast high-availability role` on each Edge to identify the two nodes participating in multicast. Also note that since unicast reachability to NSX in an active-active cluster is via ECMP, it is imperative that the northbound PIM router selects the ECMP path that matches a PIM neighbor to send PIM Join/Prune messages to NSX. In this way it will select the active Edge node which is running PIM.
- (NSX 3.2.1 and later) The NSX Edge cluster can be in active-active or active-standby mode. Note that since unicast reachability to NSX in an active-active cluster is via ECMP, it is imperative that the northbound PIM router selects the ECMP path that matches a PIM neighbor to send PIM Join/Prune messages to NSX. In this way it will select the active Edge node which is running PIM.
- East-west multicast replication: up to 4 VTEP segments for maximum replication efficiency.
- ESXi host and NSX Edge only.
- Layer 2 bridge attached to a downlink segment not supported.
- Gateway Firewall services are not supported for multicast. Distributed Firewall and Bridge Firewall are supported.
- Multi-site (NSX Federation) not supported.
- Multi-VRF not supported.

Multicast Configuration Prerequisites

Underlay network configurations:

- Acquire a multicast address range from your network administrator. This will be used to configure the Multicast Replication Range when you configure multicast on a tier-0 gateway (see Configure Multicast).

- Enable IGMP snooping on the layer 2 switches to which GENEVE participating transport nodes are attached. If IGMP snooping is enabled on layer 2, IGMP querier must be enabled on the router or layer 3 switch with connectivity to multicast enabled networks.

Multicast Configuration Steps

1 Create an IGMP profile. See Create an IGMP Profile.

2 Optionally create a PIM profile to configure a Static Rendezvous Point (RP). See Create a PIM Profile.

3 Configure a tier-0 gateway to support multicast. See Add a Tier-0 Gateway and Configure Multicast.

4 Optionally configure tier-1 gateways to support multicast. See Add a Tier-1 Gateway.

Create an IGMP Profile

Internet Group Management Protocol (IGMP) is a multicast protocol used in IPv4 networks.

Note that the IGMP snooping timeout for reports is 2 times the general query timeout. By default, the IGMP snooping timeout value is 120 seconds. On ESXi, the default IGMP snooping timeout value is 60 seconds.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Networking > Networking Profiles**.

3 Click the **Multicast Profiles** tab.

4 Click **Add IGMP Profile**.

5 Enter a profile name and the following profile details.

Option	Description
Query Interval (seconds)	Interval between general query messages. A larger value causes IGMP queries to be sent less often. Default: 30. Range: 1 - 1800.
Query Max Response Time (seconds)	Maximum allowed time before sending a response to a membership query message. Default: 10. Range: 1 - 25.
Last Member Query Interval (seconds)	Maximum amount of time between group-specific query messages, including those sent in response to leave-group messages. Default: 10. Range: 1 - 25.
Robustness Variable	Number of IGMP query messages sent. This helps alleviate the risk of loss of packets in a busy network. A larger number is recommended in a network with high traffic. Default: 2. Range: 1 - 255.

Create a PIM Profile

Protocol Independent Multicast (PIM) is a collection of multicast routing protocols for IP networks. It is not dependent on a specific unicast routing protocol and can leverage whichever unicast routing protocols are used to populate the unicast routing table.

This step is optional. It is needed only if you want to configure a Static Rendezvous Point (RP). A Rendezvous Point is a router in a multicast network domain that acts as a shared root for a multicast shared tree. If a Static RP is configured, it is preferred over the RPs that are learned from the elected Bootstrap Router (BSR).

Procedure

1. With admin privileges, log in to NSX Manager.
2. Select **Networking > Networking Profiles**.
3. Click the **Multicast Profiles** tab.
4. In the **Select Profile type** drop-down menu, select **PIM Profiles**.
5. Click **Add PIM Profile**.
6. Enter a profile name.
7. Enable or disable Bootstrap Message (BSM) processing.
8. Add one or more Static Rendezvous Point (RP) addresses.
9. Click **Save**.

About IGMP Join

When local IGMP join command for a group is configured on the uplink of a tier-0 service router (SR), IGMP group is joined on the uplink interface on the Edge Linux.

The designated router (DR) on the uplink VLAN interface will issue a PIM (*,g) join towards the Rendezvous Point (RP) upon receiving an IGMP report for the group g.

Traffic flow:

1. IGMP join is configured on the uplink of the non-active multicast node of the tier-0 SR and multicast traffic is forwarded to the non-active node from the RP. Since north-south traffic on non-active multicast node is not processed regardless of it is DR or non-DR for the uplink interface, the packet will be dropped on the non-active node. (s,g) mroute will not be present.
2. IGMP join is configured on the uplink of the active multicast node of the tier-0 SR and multicast traffic is forwarded to the active node from the RP.
    - If the node is non-DR for the uplink interface, the traffic will be dropped.

- If the node is DR for the uplink interface:
  - If IGMP join is given on the uplink interface which is also the interface to reach Source, then the incoming interface will be same as the outgoing interface for the (s,g) and the mroute will not be installed.
  - If IGMP join is given on the uplink interface which is different from the incoming interface for a given (s,g) then the (s,g) mroute will be installed.

Add an EVPN/VXLAN VNI Pool

You can create a VNI pool to be used when you configure EVPN for a tier-0 gateway. VNI pools cannot have values that overlap.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Select **Networking > Global Networking Config**.
3. Click the **EVPN/VXLAN VNI Pool** tab.
4. Click **Add EVPN/VXLAN VNI Pool**.
5. Enter a name for the pool.
6. Enter a start value.

   The value must be from 75001 to 16777215.
7. Enter an end value.

   The value must be from 75001 to 16777215.
8. Click **Save**.

Configure Global Gateway Settings

You can configure some settings that apply to all gateways.

The settings **EVPN BFD Profile** and **Enable EVPN BFD** are required by the EVPN feature to support layer-2 ECMP and BFD. For more information, see Chapter 12 Ethernet VPN (EVPN). For information about adding or editing a BFD profile, see Add a BFD Profile.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Select **Networking > Global Networking Config**.
3. In the **Global Networking Config** tab, click **Edit**.

4 (Optional) Update any of the following settings.

Option	Description
Gateway Interface MTU	The default is 1500.
Layer-3 Forwarding Mode	The options are **IPv4 Only** and **IPv4 and IPv6**. The default is **IPv4 Only**.
EVPN BFD Profile	The default is **default-external-gw-bfd-profile**.
Enable EVPN BFD	The default is **On**.

5 Click **Save**.

Add a Gateway QoS Profile

Create a QoS profile for your tier-1 gateways to define limits on the traffic rates. You can specify the permitted information rate and the burst size to set the limitations. Any traffic that does not conform to the QoS policy, is dropped. QoS profiles can be set for both ingress and egress traffic, for all traffic types (unicast, BUM, IPv4/IPv6). You can choose to create a different profile for each tier-1 gateway.

Note Gateway QoS profile is supported only on tier-1 gateways.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Networking > Networking Settings**.

3 Click the **Gateway QoS Profiles** tab.

4 Click **Add Gateway QoS Profiles**.

5 Enter a name for the profile.

6 Enter the commited bandwidth limit that you want to set for the traffic.

7 Enter the burst size. Use the following guidelines for burst size.

- *B* is the burst size in bytes.
- *R* is the committed rate (or bandwidth) in Mbps.
- *I* is the time interval in milliseconds, to refill or withdraw tokens(in bytes) from the token bucket. Use the `get dataplane` command from the NSX Edge CLI to retrieve the time interval, *Qos_wakeup_interval_ms*. The default value for *Qos_wakeup_interval_ms* is 50ms. However, this value is automatically adjusted by the dataplane based on the QoS configuration.

The constraints for burst size are:

- `B >= R * 1000,000 * I / 1000 / 8` because burst size is the maximum amount of tokens that can be refilled in each interval.

- `B >= R * 1000,000 * 1 / 1000 / 8` because the minimum value for *I* is 1 ms, taking into account dataplane CPU usage among other constraints.
- `B >= MTU of SR port` because at least the MTU-size amount of tokens need to be present in the token bucket for an MTU-size packet to pass rate-limiting check.

Since the burst size needs to satisfy all three constraints, the configured value of burst size would be:

```
Max (R * 1000,000 * I / 1000 / 8, R * 1000,000 * 1 / 1000 / 8, MTU)
```

For example, if *R* = 100 Mbps, *I* = 50 ms, and *MTU* = 1500, then

```
B >= max (100 * 1000,000 * 50 / 1000/ 8, 100 * 1000,000 * 50 / 1000/ 8, 1500) = 625000 in bytes
```

8 Click **Save**.

Add a BFD Profile

BFD (Bidirectional Forwarding Detection) is a protocol that can detect forwarding path failures. You can create a BFD profile for your Tier-0 static routes.

Note In NSX 4.0.0.1, only IPv4 is supported. Starting in NSX 4.0.1.1, both IPv4 and IPv6 are supported.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Networking** > **Networking Profiles**.

3 Click the **BFD** tab.

4 Click **Add BFD Profile**.

5 Enter a name for the profile.

6 Enter values for the heartbeat **Interval** and **Declare Dead Multiple**.

7 Click **Save**.

Add a DHCP Profile

Before you can configure DHCP on a segment, you must add a DHCP profile in your network. You can create two types of DHCP profiles: DHCP server profile and DHCP relay profile.

A DHCP profile can be used simultaneously by multiple segments and gateways in your network.

- On an overlay segment that is connected to the downlink interface of a tier-0 or a tier-1 gateway, you can attach either a DHCP server profile or a DHCP relay profile.

- On an overlay or VLAN segment that is connected to the service interface of a tier-0 or a tier-1 gateway, you can attach either a DHCP server profile or a DHCP relay profile.
- On an isolated segment (overlay or VLAN) that is not connected to a gateway, you can attach only a DHCP server profile. Isolated segment supports only a Segment DHCP server.

Add a DHCP Server Profile

You can add multiple DHCP server profiles in your network. Further, you can attach a single DHCP server profile to multiple DHCP servers.

Prerequisites

- Edge nodes are deployed in the network.
- Edge cluster is added in the network.

Procedure

1. From your browser, log in with **admin** privileges to an NSX Manager at `https://nsx-manager-ip-address`.
2. Select **Networking** > **Networking Profiles**.
3. Click the **DHCP** tab, and then click **Add DHCP Profile**.
4. Enter a unique name to identify the DHCP server profile.
5. In the **Profile Type** drop-down menu, ensure that **DHCP Server** is selected.
6. (Optional) Enter the IP address of the DHCP server in a CIDR format.

   Note A maximum of two DHCP server IP addresses are supported. You can enter one IPv4 address and one IPv6 address. For an IPv4 address, the prefix length must be <= 30, and for an IPv6 address, the prefix length must be <= 126. The DHCP server IP address must not overlap with the addresses used in DHCP ranges and DHCP static binding.

   If no server IP address is specified, 100.96.0.1/30 is autoassigned to the DHCP server.

   The server IP address cannot be any of the following:

   - Multicast IP address
   - Broadcast IP address
   - Loopback IP address
   - Unspecified IP address (address with all zeroes)

   Caution After a DHCP server profile is used in your network, preferably avoid editing the server IP addresses in the DHCP server profile. It might cause a failure while renewing or releasing the IP addresses that are leased to the DHCP clients.

7 Select an Edge cluster.

Follow these guidelines:

- If you are using a Segment DHCP server on a segment, you must select an edge cluster in the DHCP server profile. If an edge cluster is unavailable in the profile, an error message is displayed when you save the segment.
- If you are using a Gateway DHCP server on the segment, select an edge cluster either in the gateway, or DHCP server profile, or both. If an edge cluster is unavailable in either the profile or the gateway, an error message is displayed when you save the segment.

Caution After a DHCP server profile is used in your network, preferably avoid changing the edge cluster in the DHCP server profile. It might lead to a loss of existing DHCP leases that are assigned to the DHCP clients.

When a DHCP server profile is attached to a segment that uses a Segment DHCP server, the DHCP service is created in the edge cluster that you specified in the DHCP profile. However, if the segment uses a Gateway DHCP server, the edge cluster in which the DHCP service is created depends on a combination of several factors. For a detailed information about how an edge cluster is selected for DHCP service, see Scenarios: Selection of Edge Cluster for DHCP Service.

8 By default, the edges for the DHCP server are autoallocated from the edge cluster and DHCP HA is configured.

NSX selects a pair of active and standby edge nodes automatically from the available nodes in the edge cluster. If you want to allocate the edges manually from the edge cluster, go to the next step.

9 (Optional) Manually allocate the edges from the cluster.

a Turn off the **Auto Allocate Edges** toggle button.

b (Required) From the first drop-down menu, select an edge node.

c (Optional) From the second drop-down menu, select another edge node.

The first edge node becomes the active edge, and the second edge node becomes the standby edge. If second edge node is not selected, DHCP HA is not configured.

10 (Optional) Click the **Standby Relocation** toggle button to enable standby relocation. By default, this option is set to No.

Standby relocation means that if the edge node where the active or standby DHCP server is running fails, a new standby DHCP server is created on another edge node to maintain high availability. That is, if the edge node where the active DHCP server is running fails, the standby DHCP server becomes the active DHCP server and a new standby DHCP server is created on another edge node. On similar lines, if the edge node where the standby DHCP server is running fails, a new standby DHCP server replaces it.

To enable Standby Relocation, any one of these conditions must be satisfied:

- Auto Allocate Edges is set to Yes.
- Auto Allocate Edges is set to No and standby edge node is not selected.

11 (Optional) In the **Tag** drop-down menu, enter a tag name. When you are done, click **Add Item(s)**.

The maximum length of the tag name is 256 characters.

If tags exist in the inventory, the **Tag** drop-down menu displays a list of all the available tags and their scope. The list of available tags includes user-defined tags, system-defined tags, and discovered tags. You can select an existing tag from the drop-down menu and add it to the DHCP profile.

12 Click **Save**.

Add a DHCP Relay Profile

You can add a DHCP relay profile to relay the DHCP traffic to remote DHCP servers. The remote or external DHCP servers can be in any subnet, outside the SDDC, or in the physical network.

Procedure

1 From your browser, log in with **admin** privileges to an NSX Manager at `https://`*`nsx-manager-ip-address`*.

2 Select **Networking > Networking Profiles**.

3 Click the **DHCP** tab, and then click **Add DHCP Profile**.

4 Enter a unique name to identify the relay profile.

5 In the **Profile Type** drop-down menu, select **DHCP Relay**.

6 (Required) Enter the IP addresses of the remote DHCP servers.

Both DHCPv4 and DHCPv6 servers are supported. You can enter multiple IP addresses. The server IP addresses of the remote DHCP servers must not overlap with the addresses that are used in DHCP ranges and DHCP static binding.

The server IP address cannot be any of the following:

- Multicast IP address
- Broadcast IP address
- Loopback IP address
- Unspecified IP address (address with all zeroes)

7 (Optional) In the **Tag** drop-down menu, enter a tag name. When you are done, click **Add Item(s)**.

The maximum length of the tag name is 256 characters.

If tags exist in the inventory, the **Tag** drop-down menu displays a list of all the available tags and their scope. The list of available tags includes user-defined tags, system-defined tags, and discovered tags. You can select an existing tag from the drop-down menu and add it to the DHCP profile.

8 Click **Save**.

Security

16

The topics in this section cover north-south and east-west security for distributed firewall rules, identity firewall, network introspection, gateway firewall, and endpoint protection policies.

This chapter includes the following topics:

- Firewall Rule Enforcement
- Security Overview
- NSX Guest Introspection Platform
- Security Monitoring
- Security Terminology
- Identity Firewall
- Layer 7 Context Profile
- Distributed Firewall
- Gateway Firewall
- Distributed Security for vSphere Distributed Switch
- Endpoint Protection
- East-West Network Security - Chaining Third-party Services
- North-South Network Security - Inserting Third-party Service
- Network Introspection Settings
- NSX IDS/IPS and NSX Malware Prevention
- NSX Network Detection and Response
- Time-Based Firewall Policy
- Troubleshooting Firewall
- Bare Metal Server Security
- General Security Settings

Firewall Rule Enforcement

Configure east-west and north-south firewall policies under predefined categories for your environment.

Distributed Firewall (east-west) and Gateway Firewall (north-south) offer multiple sets of configurable rules divided by categories. You can configure an exclusion list that contains logical switches, logical ports, or groups, to be excluded from firewall enforcement.

Security policies are enforced as follows:

- Rules are processed in categories, left to right.
- Rules are processed in top-to-bottom ordering.
- Each packet is checked against the top rule in the rule table before moving down the subsequent rules in the table.
- The first rule in the table that matches the traffic parameters is enforced.

No subsequent rules can be enforced as the search is then terminated for that packet. Because of this behavior, it is always recommended to put the most granular policies at the top of the rule table. This ensures they will be enforced before more generic rules.

Whether an east-west or north-south firewall fails close or fails open upon failure depends on the last rule in the firewall. To ensure that a firewall fails close upon failure, configure the last rule to reject or drop all packets.

Security Overview

The Security dashboard helps you to configure features to protect your network and workloads. The **Security Overview** dashboard displays various threat detection and response features, a visual summary of the overall security configuration, and the capacity of the various objects in the NSX environment.

The information displayed on this dashboard depends on the security features that are deployed and activated in your data center.

Threat Detection & Response

This tab provides key insights about the current state of various security issues in your data center. These features help security teams understand what is happening in the network and where to focus.

Campaigns

A campaign is a set of related threat events that use specific MITRE tactics and techniques. The threat events can be mapped to MITRE ATT&CK stages to define an attack story. Campaigns can range from a single group of detection events over a short period of time to complex multi-pronged attacks over an extended amount of time. A campaign lets you view the full threat event timeline so you can respond and triage it quickly.

If the VMware NSX® Network Detection and Response™ feature is activated, this widget shows the following campaign statistics.

- The total number of campaigns that NSX Network Detection and Response has identified during the time period and that are currently active in your network.
- The total number of high impact campaigns that are in-progress during the selected time period.
- The total number of open high impact campaigns during the selected time period.
- The total number of VMs affected by the campaigns identified during the selected time period.

Click **Go to Campaigns** to see more details from the **Campaigns** page of the NSX Network Detection and Response user interface. To learn more about the NSX Network Detection and Response feature, see NSX Network Detection and Response.

IDS/IPS

- The IDS/IPS summary screen displays the following:

Entry	Description
Intrusion Events	Displays the total number of intrusion events as a clickable link, and number of intrusions that resulted in alerts or prevention.
Unique Intrusion Signatures	Displays a graph with number of intrusions detected in each severity category.
Events By Top Attack Types	Displays a graph based on attack types.

- Distributed IDS/IPS Summary

Entry	Description
Trending by Intrusion Severity	Displays a graph with the trending severity with the number of intrusion events by time.
Distribution	Displays a radar chart to show distribution based on Attack Type, Attack Target, or Severity over a period of 48 hours to 14 days.
Top VMs	Displays top VMs on which intrusion was attempted.

- Gateway IDS/IPS Summary

Entry	Description
Trending by Intrusion Severity	Displays a graph with the trending severity with the number of intrusion events by time.
Distribution	Displays a radar chart to show distribution based on Attack Type, Attack Target, or Severity over a period of 48 hours to 14 days.
Top IPs	Displays top IPs on which intrusion was attempted.

FQDN Analysis

The FQDN analysis summary screen displays:

- The total number of URLs inspected, and their severity level.
- The top URL categories that have the greatest number of inspected FQDNs.
- The highest severity URLs, with the date and the time.

URL Filtering

Select a specific gateway, or all gateways to view following information:

- Distribution of URLs by severity rating.
- Severity level of allowed URLS and displays the top five categories that have the greatest number of inspected URLs.
- Highlights the top five URL categories that have the greatest number of blocked URLs.
- Unique site distribution displays the top five sites that have the greatest number of allowed URLs. Highlights the top five sites that have the greatest number of blocked URLs.

Malicious IPs

For Distributed Firewall, you can setup Malicious IP Feed to download a list of known malicious IPs. You can block access to these IPs through firewall rules and monitor the system for any exceptions. The monitoring screen shows three charts with the following information.

- Top blocked IPs along with the total number of times the IPs are blocked.
- Top VMs accessing or accessed by malicious IPs along with the total count of malicious IPs that accessed or are accessed by the VMs.
- Top blocked categories along with the total number of times the categories are blocked.

The system also displays top 5 items of each data group.

Clicking any data point on the chart opens the Filtering and Analysis page with the detailed information about that data point. Note that the filter on the page is set to the data point you clicked. You can remove the filter and view the list of all malicious IPs.

Malware Prevention

Shows the following file events statistics for a selected time period in a graphical format:

- Total number of inspected file events, malicious file events, suspicious file events, and blocked files.
- Number of file inspections for different ranges of threat score.
- Top five recently inspected files in the data center sorted by the timestamp.
- Top five malicious files detected in the data center.

- Trend of malicious file events, suspicious file events, and suppressed file events in the data center.
- Distribution of file inspections based on the malware family to which the files belong.
- Breakdown of file inspections by the type of analysis performed (local file analysis, cloud file analysis).

Suspicious Network Activity

If VMware NSX® Intelligence™ is activated, this tab displays the following statistics (in graphical format) about suspicious or anomalous events detected during the selected time period.

- A circle shows the total number of anomalies detected during the selected time period. The circle is composed of colored segments representing the number of detected anomalous events and the MITRE adversarial tactics and technique used to detect the events.
- A list of detected suspicious events categorized in the same MITRE tactics and techniques used in their detection, and the number of times they occurred during the selected time period.
- A bar graph showing the number of anomalies detected, categorized by their severity.

Click **View All** to see more information about the detected suspicious events using the **Suspicious Traffic** page. To learn more about the NSX Suspicious Traffic feature, see the *Using and Managing VMware NSX Intelligence* documentation for version 3.2 and later at https://docs.vmware.com/en/VMware-NSX-Intelligence/index.html.

TLS Inspection

TLS inspection and decryption provides a secure way to target the influx of threats present in Enterprise web traffic. The feature uses TLS proxy to intercept encrypted traffic transparently over TLS connections and allow NSX security services such as layer 7 firewalls, IDS, and URL filtering to inspect content and enforce your security policies. You can use a wizard or manually follow the workflow to set your policy and rules.

The Security Overview dashboard shows the following TLS connection and certificate details when activated.

- The donut chart shows the TLS Connection Summary details including:
  - Bypassed due to failures
  - Decrypted
  - Connection failures
  - Bypassed due to rules
- Connections & Rules
  - Total connections

        - Open connections
        - CPS
        - Rule hits
    - The donut chart shows the Certificate Caching details including:
        - Cache hits
        - Cached certificates
        - Cache misses
    - Traffic
        - Throughput details including Client to server and Server to Client
        - Total traffic details including Client to server and Server to Client

Configuration

The **Configuration** tab provides a quick summary view with clickable links with the number of:

- Firewall Policies
- Endpoint Policies
- IDS/IPS Policies
- Malware Prevention Policies
- Network Introspection policies
- TLS Inspection Policies

This page also provides detailed views of security settings for:

Gateway Firewall widget

Highlights gateway firewall security settings. Click the links to view the gateways on which the following security features are activated:

- IDS/IPS
- Malware Prevention
- TLS Inspection

To view the gateways with these security features, at least one of the above security features must be deployed in your data center.

Distributed Firewall widget

Highlights the total distributed firewall policies using graphics. Click to view details such as policy groupings, top services consumed by East-West security policies as well as their actions (allow, drop, and reject), and total distributed firewall rules.

Endpoint Protection widget

Shows a summary of the configuration of endpoint protection for virtual machines. You can view VM distribution by service profile, components having issues, and configured VMs running file introspection.

Identity Firewall User Sessions widget

Displays the number of IDFW active user sessions.

Malware Prevention widget

This UI widget shows issues when any of the components for the NSX Distributed Malware Prevention service is down or not working.

For example:

- The Bar chart shows an issue when the Security Hub on the NSX Malware Prevention service virtual machine (SVM) is down. Point to the bar to view the following details:
  - Number of NSX Malware Prevention SVMs that are impacted.
  - Number of workload VMs on the host that have lost malware security protection due to the Security Hub going down.
- The Donut chart shows the following details:
  - Number of workload VMs where the NSX File Introspection driver is running.
  - Number of workload VMs where the NSX File Introspection driver is not running.

  For both these metrics, only the workload VMs on the host clusters that are activated for NSX Distributed Malware Prevention are considered.

Capacity

Capacity information is available only in the Manager mode of the NSX Manager UI. The information displayed on this dashboard depends on the security features that are deployed and activated in your data center. It shows details for:

- Introspection Rules N-S Tier-1
- Active Directory Domains (Identity Firewall)
- Service Chains
- Introspection Policies E-W
- Saved Firewall Rules Configuration
- Introspection Policies N-S Tier-0

- Introspection Service Paths
- System-wide Firewall Rules
- Introspection Rules N-S Tier-0
- Introspection Rules E-W
- System-wide Endpoint Protection Enabled Virtual Machines
- Introspection Policies N-S Tier-1
- Distributed Firewall Sections
- System-wide Firewall Sections
- Active Directory Domains (Identity Firewall)
- System-wide Endpoint Protection Enabled Hosts
- Distributed Firewall Rules

NSX Guest Introspection Platform

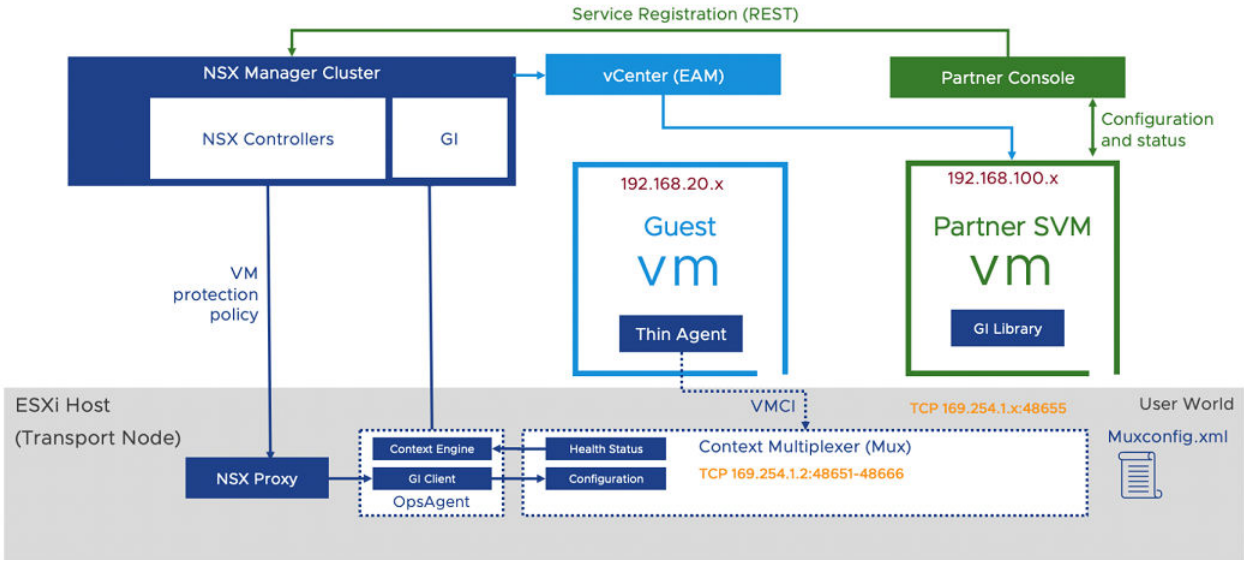
NSX Malware Prevention, Endpoint Protection, IDFW and NSX Intelligence features use the functionality provided by NSX Guest Introspection Platform.

NSX Guest Introspection Platform Architecture

This topic helps describes the key concepts and architecture of service insertion and NSX Guest Introspection Platform components (guest introspection).

Figure 16-1. Endpoint Protection Architecture (NSX Guest Introspection Platform Use Case)

Figure 16-2.



Note In the diagram, blocks in blue are NSX Guest Introspection Platform components and blocks in green are Endpoint Protection Consumer components.

Key Concepts:

- NSX Manager: It is the management plane appliance for NSX that provides API and graphical user interface to customers and partners for configuration of Network and Security policies. For guest introspection, the NSX Manager also provides API and GUI to deploy and manage partner appliances.
- Guest Introspection SDK: VMware provided library consumed by the security vendor.
- Service VM: It is the security vendor provided VM that consumes the guest introspection SDK provided by VMware. It contains the logic to scan file, process, network and system events to detect virus or malware on the guest. After scanning a request, it sends back a verdict or notification about the action taken by the guest VM on the request.
- Guest Introspection host agent (Context Multiplexer): It processes configuration of endpoint protection policies. It also multiplexes and forwards messages from protected VMs to the Service VM and back from Service VM to protected VMs. It reports the health status of the guest introspection platform and maintains records of the Service VM configuration in the `muxconfig.xml` file.
- Ops agent (Context engine and Guest Introspection client): It forwards the guest introspection configuration to the guest introspection host agent (Context Multiplexer). It also relays the health status of the solution to NSX Manager.
- EAM: NSX Manager uses the ESXi agent manager to deploy a partner Service VM on every host on the cluster configured for protection.

- Thin agent: It is the file and/or network introspection agent running inside the guest VMs. It intercepts file and network activities or events that are forwarded to the Service VM through the host agent. On Windows, this agent is part of VMware Tools. On Linux, this agent exists outside of VMware Tools. It provides in-guest context to NSX and can also be used to replace the traditional agent provided by antivirus or antimalware security vendors. It is a generic and lightweight agent that facilitates offloading files and processes for scanning to the Service VM provided by the vendor.

The benefits of using the guest introspection platform to protect guest VM endpoints:

- Reduced consumption of compute resources: Guest introspection offloads signatures and security scanning logic from each endpoint on a host to a third-party partner Service VM on the host. As scanning happens only on the Service VM, there is no need to spend compute resources on guest VMs to run scans.
- Better management: As scans are offloaded to a Service VM, signatures need to be updated to only one object per host. Such a mechanism works better than agent-based solution where same signatures need updates on all guest VMs.
- Continuous antivirus and antimalware protection: As the Service VM runs continuously, a guest VM is not mandated to run the latest signatures. For example, a snapshot VM might run some older version of the signature making it vulnerable in the traditional way of protecting endpoints. With the guest introspection platform, the Service VM is continuously running the latest and malware signatures thereby ensuring that any newly added VM is also protected with the latest signatures.
- Offloaded signatures to a Service VM: database lifecycle is outside of guest VM lifecycle and so the Service VM is not affected by guest VM outages.

NSX Guest Introspection Platform Use Cases

NSX Guest Introspection Platform provide guest introspection services to a few of the NSX security features.

Supported use cases of NSX Guest Introspection Platform are:

- NSX Endpoint Protection: See NSX Guest Introspection Platform Architecture.
- NSX Malware Prevention: See Overview of NSX IDS/IPS and NSX Malware Prevention.
- Identity Firewall: See Identity Firewall.
- NSX Intelligence: See the *Using and Managing VMware NSX Intelligence* guide.

NSX Endpoint Protection and NSX Malware Prevention Use Case

In a virtual environment, the NSX Guest Introspection Platform enables provision of agentless security solutions for guest VMs.

As an NSX administrator, you implement an antivirus and antimalware solution that is deployed as a Service Virtual Machine (Service VM, or SVM) to monitor a file, network, or process activity on a guest VM. Whenever a file is accessed, such as a file open attempt, the antimalware Service VM is notified of the event. The Service VM then determines how to respond to the event. For example, to inspect the file for signatures.

- If the Service VM determines that the file contains no malware, then it allows the file open operation to succeed.
- If the Service VM detects a in the file, it requests the Thin Agent on the guest VM to act in one of the following ways:
  - Delete the infected file or deny access to the file.
  - Infected VMs can be assigned a tag by NSX. Moreover, you can define a rule that automatically moves such tagged guest VMs to a security group that quarantines the infected VM for additional scan and isolation from the network until the infection is completely removed.

Installing Host Components

Before you install the guest components, ensure the host components are installed on the system.

Install the following host components:

- MUX and Ops-agent: Prepare the host as a transport node. See *NSX Installation Guide*.
- NSX Endpoint Protection installation:
  - To register and configure endpoint protection services with NSX, see the third-party partner documentation.
  - To deploy and consume NSX Endpoint Protection rules, see Endpoint Protection.
- NSX Malware Prevention: To enable and configure NSX Malware Prevention, see NSX IDS/IPS and NSX Malware Prevention.
- Guest Introspection SDK: To configure Guest Introspection SDK, see Guest Introspection SDK documentation.

Installing Guest Components

Install guest components on Linux and Windows virtual machines. Guest components include installing file introspection and network introspection drivers.

Install the Guest Introspection Thin Agent for Anti-virus on Linux Virtual Machines

Guest Introspection supports File Introspection in Linux for anti-virus only. To protect Linux VMs using a Guest Introspection security solution, you must install the Guest Introspection thin agent.

The Linux thin agent is available as part of the operating system specific packages (OSPs). The packages are hosted on VMware packages portal. Enterprise or Security Administrator (non-NSX Administrator) can install the agent on guest VMs outside of NSX.

Installing VMware Tools is not required.

Based on your Linux operating system, perform the following steps with root privilege:

Prerequisites

- Ensure that the guest virtual machine has a supported version of Linux installed:
  - Red Hat Enterprise Linux (RHEL) 7.6, 7.7, 8.2 (64 bit) GA
  - SUSE Linux Enterprise Server (SLES) 12 SP3+, 15 SP1 (64 bit) GA
  - Ubuntu 16.04.5, 16.04.6, 18.04, 20.04 (64 bit) GA
  - CentOS 7.6, 7.7, 8.2 (64 bit) GA
- Verify GLib 2.0 is installed on the Linux VM.

Procedure

1 For Ubuntu systems

a Obtain and import the VMware packaging public keys using the following commands.

```
curl -O https://packages.vmware.com/packages/nsx-gi/keys/VMWARE-PACKAGING-NSX-GI-GPG-RSA-KEY.pub
apt-key add VMWARE-PACKAGING-NSX-GI-GPG-RSA-KEY.pub
```

b Create a new file named `vmware.list` file under `/etc/apt/sources.list.d`

c Edit the file with the following content:

For Ubuntu 16.04

```
deb [arch=amd64] https://packages.vmware.com/packages/nsx-gi/latest/ubuntu/ xenial main
```

For Ubuntu 18.04

```
deb [arch=amd64] https://packages.vmware.com/packages/nsx-gi/latest/ubuntu/ bionic main
```

For Ubuntu 20.04

```
deb [arch=amd64] https://packages.vmware.com/packages/nsx-gi/latest/ubuntu/ focal main
```

d Install the package.

```
apt-get update
apt-get install vmware-nsx-gi-file
```

2 For RHEL7 systems

a Obtain and import the VMware packaging public keys using the following commands.

```
curl -O https://packages.vmware.com/packages/nsx-gi/keys/VMWARE-PACKAGING-NSX-GI-GPG-RSA-KEY.pub
rpm --import VMWARE-PACKAGING-NSX-GI-GPG-RSA-KEY.pub
```

b Create a new file named `vmware.repo` file under `/etc/yum.repos.d`.

c Edit the file with the following content:

```
[vmware]
name = VMware
baseurl = https://packages.vmware.com/packages/nsx-gi/latest/rhel/x86_64
enabled = 1
gpgcheck = 1
metadata_expire = 86400
ui_repoid_vars = basearch
```

d Install the package.

```
yum install vmware-nsx-gi-file
```

3 For SLES systems

a Obtain and import the VMware packaging public keys using the following commands.

```
curl -O https://packages.vmware.com/packages/nsx-gi/keys/VMWARE-PACKAGING-NSX-GI-GPG-RSA-KEY.pub
rpm --import VMWARE-PACKAGING-NSX-GI-GPG-RSA-KEY.pub
```

b Add the following repository:

```
zypper ar -f "https://packages.vmware.com/packages/nsx-gi/latest/sles/x86_64/" VMware
```

c Install the package.

```
zypper install vmware-nsx-gi-file
```

4 For CentOS systems

a Obtain and import the VMware packaging public keys using the following commands.

```
curl -O https://packages.vmware.com/packages/nsx-gi/keys/VMWARE-PACKAGING-NSX-GI-GPG-RSA-KEY.pub
rpm --import VMWARE-PACKAGING-NSX-GI-GPG-RSA-KEY.pub
```

b Create a new file named `vmware.repo` file under `/etc/yum.repos.d`.

c Edit the file with the following content:

```
[vmware]
name = VMware
baseurl = https://packages.vmware.com/packages/nsx-gi/latest/centos/x86_64
enabled = 1
gpgcheck = 1
metadata_expire = 86400
ui_repoid_vars = basearch
```

d Install the package.

```
zypper install vmware-nsx-gi-file
```

What to do next

Verify whether the thin agent is running using the `service vsepd status` or `systemctl status vsepd`command with the administrative privileges. The status must be running.

Install the Linux Thin Agent for Network Introspection

Install the Linux thin agent to introspect network traffic. The network introspection driver will be used by NSX Malware Prevention functionality to introspect traffic for any malware.

Important To protect guest VMs against antivirus, you do not need to install the Linux thin agent for network introspection.

The Linux thin agent driver that is used to introspect network traffic depends on an open-source driver.

Prerequisites

Install the following packages:

- glib2
- libnetfilter-conntrack3/ libnetfilter-conntrack
- libnetfilter-queue1/ libnetfilter-queue
- iptables

Procedure

1 (Ubuntu)

a To install the open-source driver provided by guest introspection, add following in `/etc/apt/sources.list` as the base URL for your operating system.

```
deb [arch=amd64] https://packages.vmware.com/guest-introspection-for-vmware-nsx/latest/
ubuntu xenial main
```

b Import the VMware packaging key.

```
curl -O https://packages.vmware.com/packages/nsx-gi/keys/VMWARE-PACKAGING-NSX-GI-GPG-
RSA-KEY.pub
apt-key add VMWARE-PACKAGING-NSX-GI-GPG-RSA-KEY.pub
```

c Update the repository and install the open-source driver.

```
apt-get install guest-introspection-for-vmware-nsx
```

d To install the Linux thin agent that is used to introspect file and or network traffic.

- To install file and network introspection packages, select `vmware-nsx-gi` package in step g.
- To install network introspection packages, select the `vmware-nsx-gi-net` package in step g.

e Create `/etc/apt/source.list.d/gi.list` and add following URL as the base URL for your operating system depending upon the distro.

```
deb [arch=amd64] http://packages.vmware.com/packages/nsx-gi/latest/ubuntu xenial main
```

Or

```
deb [arch=amd64] http://packages.vmware.com/packages/nsx-gi/latest/ubuntu focal main
```

f Import the VMware packaging key.

```
curl -O https://packages.vmware.com/packages/nsx-gi/keys/VMWARE-PACKAGING-NSX-GI-GPG-
RSA-KEY.pub
apt-key add VMWARE-PACKAGING-NSX-GI-GPG-RSA-KEY.pub
```

g Install one of the following drivers.

```
apt-get install vmware-nsx-gi
apt-get install vmware-nsx-gi-net
```

2 (RHEL)

a To install the open-source driver provided by guest introspection, create `/etc/yum.repos.d/gi.repo` file and add following in the file.

```
[nsx-gi]
name=nsx-gi
baseurl=https://packages.vmware.com/guest-introspection-for-vmware-nsx/1.2.0.0/rhel/x86_64/
enabled=1
gpgcheck=1
```

b Import the VMware packaging key.

```
curl -O https://packages.vmware.com/packages/nsx-gi/keys/VMWARE-PACKAGING-NSX-GI-GPG-RSA-KEY.pub
rpm --import VMWARE-PACKAGING-NSX-GI-GPG-RSA-KEY.pub
```

c Update the repository and install the open-source driver.

```
yum install Guest-Introspection-for-VMware-NSX
```

d To install the Linux thin agent that is used to introspect file and or network traffic.

- To install file and network introspection packages, select `vmware-nsx-gi` package in step g.
- To install network introspection packages, select the `vmware-nsx-gi-net` package in step g.

e Create `/etc/yum.repos.d/vm.repo` file and add following:

```
[vmware]
name = VMware
baseurl = http://packages.vmware.com/packages/nsx-gi/latest/rhel/x86_64
enabled = 1
gpgcheck = 1
metadata_expire = 86400
ui_repoid_vars = basearch
```

f Import the VMware packaging key.

```
https://packages.vmware.com/packages/nsx-gi/keys/VMWARE-PACKAGING-NSX-GI-GPG-RSA-KEY.pub
```

g Install drivers.

```
yum install vmware-nsx-gi
yum install vmware-nsx-gi-net
```

3 (SLES)

a To install the open-source driver provided by guest introspection, create `/etc/yum.repos.d/gi.repo` file and add following in the file.

```
[nsx-gi]
name=nsx-gi
baseurl=https://packages.vmware.com/guest-introspection-for-vmware-nsx/1.2.0.0/sles/x86_64/
enabled=1
gpgcheck=1
```

b Import the VMware packaging key.

```
curl -O https://packages.vmware.com/packages/nsx-gi/keys/VMWARE-PACKAGING-NSX-GI-GPG-RSA-KEY.pub
rpm --import VMWARE-PACKAGING-NSX-GI-GPG-RSA-KEY.pub
```

c Update the repository and install the open-source driver.

```
zypper install Guest-Introspection-for-VMware-NSX
```

d To install the Linux thin agent that is used to introspect file and/or network traffic.

- To install file and network introspection packages, select `vmware-nsx-gi` package in step g.
- To install network introspection packages, select the `vmware-nsx-gi-net` package in step g.

e Create `/etc/yum.repos.d/vm.repo` file and add following:

```
[vmware]
name = VMware
baseurl = http://packages.vmware.com/packages/nsx-gi/latest/sles/x86_64
enabled = 1
gpgcheck = 1
metadata_expire = 86400
ui_repoid_vars = basearch
```

f Import the VMware packaging key.

```
https://packages.vmware.com/packages/nsx-gi/keys/VMWARE-PACKAGING-NSX-GI-GPG-RSA-KEY.pub
```

g Install drivers.

```
zypper install vmware-nsx-gi
zypper install vmware-nsx-gi-net
```

4 (CentOS)

a To install the open-source driver provided by guest introspection, create `/etc/yum.repos.d/gi.repo` file and add following in the file.

```
[nsx-gi]
name=nsx-gi
baseurl=https://packages.vmware.com/guest-introspection-for-vmware-nsx/1.2.0.0/centos/x86_64/
enabled=1
gpgcheck=1
```

b Import the VMware packaging key.

```
curl -O https://packages.vmware.com/packages/nsx-gi/keys/VMWARE-PACKAGING-NSX-GI-GPG-RSA-KEY.pub
rpm --import VMWARE-PACKAGING-NSX-GI-GPG-RSA-KEY.pub
```

c Update the repository and install the open-source driver.

```
yum install Guest-Introspection-for-VMware-NSX
```

d To install the Linux thin agent that is used to introspect file and or network traffic.

- To install file and network introspection packages, select `vmware-nsx-gi` package in step g.
- To install network introspection packages, select the `vmware-nsx-gi-net` package in step g.

e Create `/etc/yum.repos.d/vm.repo` file and add following:

```
[vmware]
name = VMware
baseurl = http://packages.vmware.com/packages/nsx-gi/latest/centos/x86_64
enabled = 1
gpgcheck = 1
metadata_expire = 86400
ui_repoid_vars = basearch
```

f Import the VMware packaging key.

```
https://packages.vmware.com/packages/nsx-gi/keys/VMWARE-PACKAGING-NSX-GI-GPG-RSA-KEY.pub
```

g Install drivers.

```
yum install vmware-nsx-gi
yum install vmware-nsx-gi-net
```

Install the Guest Introspection Thin Agent on Windows Virtual Machines for Anti-virus

To protect VMs from anti-virus using a Guest Introspection security solution, you must install Guest Introspection thin agent, also called Guest Introspection drivers, on the VM. Guest Introspection drivers are included with VMware Tools for Windows, but are not part of the default installation. To install Guest Introspection on a Windows VM, you must perform a custom install and select the drivers or run complete install.

Windows virtual machines with the Guest Introspection drivers installed are automatically protected whenever they are started up on an ESXi host that has the security solution installed and VM protection policies configured. Protected virtual machines retain the security protection through shut downs and restarts, and even after a vMotion move to another ESXi host with the security solution installed.

- If you are using vSphere 6.0, see these instructions for installing VMware Tools, see Manually Install or Upgrade VMware Tools in a Windows Virtual Machine.
- If you are using vSphere 6.5, see these instructions for installing VMware Tools, see Install VMware Tools on vSphere 6.5.

Prerequisites

Ensure that the guest virtual machine has a supported version of Windows installed.

- To know the which versions of Windows operating systems are supported on a particular ESXi host version, see the VMware Compatibility Guide. In the Windows Compatibility Guide, narrow your search to Guest OS: In the **What are you looking for** field, select **`Guest OS`** and in the **Product Name** field, select **`ESXi`** and click **Update and View Results**.
- To know the guest operating systems supported by VMware Tools, go to VMware Tools Documentation and see the Release Notes of VMware Tools version you want to install. For example, see the VMware Tools 12.1.5 Release Notes to know the supported Windows operating systems.

Procedure

1. Start the VMware Tools installation, following the instructions for your version of vSphere. Select **Custom install**.
2. Expand the VMCI Driver section.

   The options available vary depending on the version of VMware Tools.
3. Select the driver to be installed on the VM.

Driver	Description
NSX File Introspection Driver	Select NSX File Introspection Driver to install vsepflt.

4. In the drop-down menu next to the drivers you want to add, select This feature is installed on the local hard drive.

5 Follow the remaining steps in the procedure.

What to do next

Verify whether the thin agent is running using the `fltmc` command with the administrative privileges. The Filter Name column in the output lists the thin agent with an entry `vsepflt`.

Install the Guest Introspection Thin Agent on Windows Virtual Machines for Network Introspection

To protect VMs using a Guest Introspection security solution, you must install Guest Introspection thin agent, also called Guest Introspection drivers, on the VM. Guest Introspection drivers are included with VMware Tools for Windows, but are not part of the default installation. To install Guest Introspection on a Windows VM, you must perform a custom install and select the drivers or run complete install.

Windows virtual machines with the Guest Introspection drivers installed are automatically protected whenever they are started up on an ESXi host that has the security solution installed and VM protection policies configured. Protected virtual machines retain the security protection through shutdowns and restarts, and even after a vMotion move to another ESXi host with the security solution installed.

- If you are using vSphere 6.0, see these instructions for installing VMware Tools, see Manually Install or Upgrade VMware Tools in a Windows Virtual Machine.
- If you are using vSphere 6.5, see these instructions for installing VMware Tools: Install VMware Tools in vSphere 6.5.

Prerequisites

Ensure that the guest virtual machine has a supported version of Windows installed.

- To know the which versions of Windows operating systems are supported on a particular ESXi host version, see the VMware Compatibility Guide. In the Windows Compatibility Guide, narrow your search to Guest OS: In the **What are you looking for** field, select **`Guest OS`** and in the **Product Name** field, select **`ESXi`** and click **Update and View Results**.
- To know the guest operating systems supported by VMware Tools, go to VMware Tools Documentation and see the Release Notes of VMware Tools version you want to install. For example, see the VMware Tools 12.1.5 Release Notes to know the supported Windows operating systems.

Procedure

1 Start the VMware Tools installation, following the instructions for your version of vSphere. Select **Custom install**.

2 Expand the VMCI Driver section.

The options available vary depending on the version of VMware Tools.

3 Select the driver to be installed on the VM.

Driver	Description
vShield Endpoint Drivers	Installs Network Introspection (`vnetflt`) driver.
Guest Introspection Drivers	Installs Network Introspection (`vnetflt`) driver.
NSX Network Introspection Driver	Select NSX Network Introspection Driver to install `vnetflt` (`vnetWFP` on Windows 10 or later). **Note** Select NSX Network Introspection Driver only if you are using the Identity Firewall or Endpoint Monitoring features.

4 In the drop-down menu next to the drivers you want to add, select This feature is installed on the local hard drive.

5 Follow the remaining steps in the procedure.

What to do next

Verify whether the thin agent is running using the `sc query vnetwfp` command with the administrative privileges. The Filter Name column in the output lists the thin agent with an entry `vnetwfp`.

Supported File Systems for Guest VMs

The NSX Guest Introspection Platform SDK supports monitoring guest VMs that use the following file systems:

Windows

- NTFS
- FAT
- CDFS
- UDFS
- LANMAN (LanMan redirector)
- MUP (MUP redirector)

Linux

- EXT2
- EXT3
- EXT4
- NFS
- NFS4
- XFS

- BTRFS
- VFAT
- ISO9660
- CIFS

Logging and Troubleshooting Host Components

Troubleshoot issues related to host components such as MUX agent, EPSeclib, context engine and so on.

NSX Guest Introspection Platform Host Agent (MUX) Logs

If virtual machines on an ESXi host are not working with NSX Malware Prevention, NSX Endpoint Protection, IDFW or NSX Intelligence or if there are alarms on a host regarding communication to the SVA, then NSX Guest Introspection Platform on the ESXi host might face some issues.

Log Path and Sample Message

MUX Log path
/var/log/syslog

NSX Guest Introspection Platform host agent (MUX) messages follow the format of <timestamp>ContextMUX<[ThreadID]>: <message>

For example:

```
2022-08-25T11:42:02Z ContextMux[54425342]: [INFO] (EPSEC) [54425342] NSX Context Multiplexor
17883598. [3.1.2.0.0]
2022-08-25T11:42:02Z ContextMux[54425342]: [INFO] (EPSEC) [54425342] VMkernel localhost 6.7.0
#1 SMP Release build-13006603 Mar 26 2019 13:38:24 x86_64
2022-08-25T11:42:02Z ContextMux[54425342]: [WARNING] (EPSEC) [54425345] Not sending events.
MuxHandler has not been registered.
```

In the above example

- [ERROR] is the type of message. Other types can be [DEBUG], [INFO]
- (EPSEC) represents that the messages are specific to Endpoint Security

Enabling and Viewing Log Files

To view the version of the NSX Guest Introspection Platform host agent (MUX) VIB installed on the host, run the `#esxcli software vib list | grep nsx-context-mux` command.

To turn on full logging, perform these steps on the ESXi host command shell:

1 Set the NSX Guest Introspection Platform host agent (MUX) logging level.

```
~ # nsxcli

localhost> set service nsx-context-mux logging-level debug
```

`off` - The highest possible rank and is intended to turn off logging

`fatal` - Designates very severe events that may cause application to abort

`error` - Designates error events that may allow application to continue running

`warn` - Designates potentially harmful situations

`info` - Designates coarse-grained informational messages

`debug` - Designates fine-grained informational messages for debugging purposes

`trace` - Designates finer-grained informational events than debug

2 View the NSX Guest Introspection Platform host agent (MUX) log messages in the `/var/log/syslog.log` file on the ESXi host. Check that the entries corresponding to the global solutions, solution ID, and port number are specified correctly.

For more details, refer to NSX Guest Introspection Platform Host Agent (MUX) Logs.

Example: Sample muxconfig.xml File

```
<?xml version="1.0" encoding="UTF-8"?>
<EndpointConfig>
 <hostMoid>702e86b6-09ad-4dc7-9473-fbe2ba01f7f7</hostMoid>
 <InstalledSolutions>
  <Solution>
   <id>100</id>
   <listenOn>unixdom</listenOn>
   <port>48656</port>
   <uuid/>
   <vmxPath/>
  </Solution>
  <Solution>
   <id>7498352642083520512</id>
   <ipAddress>xxx.xxx.x.xx</ipAddress>
   <listenOn>ip</listenOn>
   <port>48651</port>
   <uuid/>
   <vmxPath/>
  </Solution>
 </InstalledSolutions>
 <DefaultSolutions/>
 <GlobalSolutions>
  <solution>
   <id>100</id>
   <tag/>
   <order>0</order>
  </solution>
 </GlobalSolutions>
 <VmConfig>
  <uuid>50259775-1c32-a940-de62-ed537ba0d0bf</uuid>
  <solution>
   <id>7498352642083520512</id>
   <tag>Gold</tag>
   <order>7498352642083520512</order>
```

```
    <action/>
    <file_types/>
   </solution>
  </VmConfig>
 </EndpointConfig>
```

Troubleshooting NSX Guest Introspection Platform Host Agent (MUX)

If all virtual machines on an ESXi host are not working with NSX Guest Introspection Platform or if there are alarms on a particular host regarding communication to the EPP SVA, then it could be a problem with the NSX Guest Introspection Platform on the ESXi host.

To enable debugs based on context (file, process, network, or system), see Troubleshooting the Thin Agent on Linux.

NSX Guest Introspection Platform Host Agent (MUX)

1 Check to see if the service is running on the ESXi host by running the *# /etc/init.d/nsx-context-mux status* command:

For example:

```
# /etc/init.d/nsx-context-mux  status
         nsx-context-mux  is running
```

2 Run the `ps -c |grep Mux` command to find the NSX Guest Introspection Platform host agent (MUX) processes that are currently running.

For example:

```
[root@localhost:~] ps -c | grep mux
54425332 54425332 sh                     /bin/sh /sbin/watchdog.sh -s nsx-context-mux -q
100 -t 1000000 /usr/lib/vmware/mux/bin/nsx-context-mux 900 -c 910 -u mux_user
54425342 54425342 nsx-context-mux              /usr/lib/vmware/mux/bin/nsx-context-mux 900
-c 910 -u mux_user
54425345 54425342 nsx-context-mux              /usr/lib/vmware/mux/bin/nsx-context-mux 900
-c 910 -u mux_user
54425346 54425342 nsx-context-mux              /usr/lib/vmware/mux/bin/nsx-context-mux 900
-c 910 -u mux_user
54425347 54425342 nsx-context-mux              /usr/lib/vmware/mux/bin/nsx-context-mux 900
-c 910 -u mux_user
54425348 54425342 nsx-context-mux              /usr/lib/vmware/mux/bin/nsx-context-mux 900
-c 910 -u mux_user
```

3 If the service is not running, you can restart it with these commands:`/etc/init.d/nsx-context-mux start`.

4 If you see that the service is not running, restart it or start it with this command: `/etc/init.d/nsx-context-mux start`.

Note that it is safe to restart this service during production hours as it does not have any great impact, and restarts in a couple of seconds.

5 To stop the MUX process, run `/etc/init.d/nsx-context-mux stop`.

6 To check the logging level currently in use, run `get service nsx-context-mux logging-level`.

7 To get a better idea of what the NSX Guest Introspection Platform is doing or check the communication status, you can check the logs on the ESXi host. NSX Guest Introspection Platform logs are written to the host `/var/log/syslog` file. This is also included in the ESXi host support logs.

   For more information, see Collecting diagnostic information for ESX/ESXi hosts and vCenter Server using the vSphere Web Client (2032892)

8 The default logging option for NSX Guest Introspection Platform is info and can be raised to debug to gather more information.

9 If you run into issues with the VIB installation, check the /var/log/esxupdate.log file on the ESXi host. This log provides accurate information to know the reason the package was not successfully installed. This is a common issue for NSX Guest Introspection Platform installation issues. For more information, see Installing NSX Guest Introspection services (ESX GI Module VIB) on the ESXi host fails in VMware NSX for vSphere 6.x (2135278).

10 To check for a corrupt ESXi image look for a message similar to this:

```
esxupdate: esxupdate: ERROR: Installation Error:
 (None, 'No image profile is found on the host or image profile is empty.
 An image profile is required to install or remove VIBs. To install an image profile,
 use the esxcli image profile install command.')
```

11 To verify that the image is corrupted run the command *cd /vmfs/volumes* on the ESXi host.

   a Search for the imgdb.tgz file by running this command: *find * | grep imgdb.tgz*.

      This command normally results in two matches. For example:

      0ca01e7f-cc1ea1af-bda0-1fe646c5ceea/imgdb.tgz or *edbf587b-da2add08-3185-3113649d5262/imgdb.tgz*

   b On each match, run this command: *ls -l match_result*

      For example:

```
> ls -l 0ca01e7f-cc1ea1af-bda0-1fe646c5ceea/imgdb.tgz -rwx------
 1 root root  26393 Jul 20 19:28 0ca01e7f-cc1ea1af-bda0-1fe646c5ceea/imgdb.tgz
   > ls -l edbf587b-da2add08-3185-3113649d5262/imgdb.tgz -rwx------
  1 root root   93 Jul 19 17:32 edbf587b-da2add08-3185-3113649d5262/imgdb.tgz
```

      The default size for the imgdb.tgz file is far greater than the other file or if one of the files is only a couple of bytes, it indicates that the file is corrupt. The only supported way to resolve this is to re-install ESXi for that particular ESXi host.

NSX Guest Introspection Platform SDK EPSecLib

The EPSecLib receives events from the ESXi host NSX Guest Introspection Platform Host Agent (MUX).

Log Path and Sample Message

EPSecLib Log Path
/var/log/syslog

EPSecLib messages follow the format of <timestamp> <VM Name><Process Name><[PID]>: <message>

In the following example [ERROR] is the type of message and (EPSEC) represents the messages that are specific to any functionality that uses NSX Guest Introspection Platform.

For example:

```
Oct 17 14:26:00 endpoint-virtual-machine EPSecTester[7203]: [NOTICE] (EPSEC)
 [7203] Initializing EPSec library build: build-00000

Oct 17 14:37:41 endpoint-virtual-machine EPSecSample: [ERROR] (EPSEC) [7533] Event
terminated reading file. Ex: VFileGuestEventTerminated@tid=7533: Event id: 3554.
```

Collecting Logs

To enable debug logging for the EPSec library, which is a component inside any service that uses NSX Guest Introspection Platform:

1. Work with the anti-virus or NSX Malware Prevention security vendor to enable console or SSH access to the SVM. Follow partner provided instructions to enable console or SSH access.
2. Log in to the EPP or NSX Malware Prevention SVM by obtaining the console password from NSX Manager.
3. Create `/etc/epseclib.conf` file and add:

   `ENABLE_DEBUG=TRUE`

   `ENABLE_SUPPORT=TRUE`

   The debug logs can be found in (RHEL/SLES/CentOS) `/var/log/messages` or (Ubuntu) `/var/log/syslog` . Because the debug setting can flood the `/var/log` file, disable the debug mode as soon as you have collected all the required information.
4. Change permissions by running the `chmod 644 /etc/epseclib.conf` command.
5. Work with the anti-virus or NSX Malware Prevention partner to extract logs generated for the SVM.
6. For NSX Malware Prevention, configure the Security Hub VM to enable EPSecLib.

Logging and Troubleshooting Guest Components

Enable logging and troubleshoot issues related to the NSX Guest Introspection Platform thin agent.

Troubleshooting the Thin Agent on Linux

The Guest Introspection thin agent is installed with VMware Tools™ on each guest virtual machine.

Troubleshooting the Thin Agent on Linux

If a virtual machine is slow in reading and writing operations, and unzipping or saving files then there might be problems with the thin agent.

1. Check the compatibility of all the components involved. You need the build numbers for ESXi, vCenter Server, NSX Manager, and the Security solution you have selected (for example, Trend Micro, McAfee, Kaspersky, or Symantec). After this data has been collected, compare the compatibility of the vSphere components. For more information, see the VMware Product Interoperability Matrices.
2. Ensure that File Introspection is installed on the system.
3. Verify that the thin agent is running by with the `systemctl status vsepd.service` command.
4. If you suspect that the thin agent is causing a performance problem with the system, stop the service by running the `service vsepd stop` command.
5. Then perform a test to get a baseline. You can then start the vsep service and perform another test by running the `service vsepd start` command.
6. For deployments consuming network events, `vmw_conn_notify` also needs to be checked by running the `systemctcl status vmw_conn_notifyd.service`.
7. Enable debugging for the Linux thin agent:
   a. Run `/etc/vsep/vsep refresh-logging`.
   b. Usage: `/etc/vsep/vsepd refresh-logging <dest> <<level> sub-component-name>`

      Where, `<dest>`: [1-2] 1 - log to the VM and 2 - log to the ESX host.

      `<level>`: **`[1-7]`** where **`4`** is for logging level `INFO`, **`7`** is for logging level `DEBUG`.

      `<sub-component-name>`: one or more of transport, timer, file, network, process, system

      When logging to host is enabled, logs are stored in `vmware.log` of respective `vmfs` directory of VMs on ESXi hosts.

      Note Enabling full logging might result in heavy log activity flooding the `vmware.log` file. Disable full logging as soon as possible.

Enable debugs based on context (file, process, network or system)

Enhanced logging support allowed thin agent to log module debug level information of specific feature/functionality to vmware.log on host or syslog in the VM.

One needs to restart Thin Agent service in case debug logs are not generated in respective files. Note that logging to `vmware.log` on host may get throttled if logging is heavy. The `refresh-logging` input parameter has been added to `/etc/vsep/vsepd`. Its usage can be displayed by running:

Debugging:

```
# /etc/vsep/vsepd refresh-logging
```

Usage: `/etc/vsep/vsepd refresh-logging <dest> <<level> sub-component-name>`

where, `<dest>`: **`[1-2]`**: **1** is log to the VM and **2** is log to the ESX host. When logging to VM is enabled, the logs will be stored at the following location based on the Linux distribution software.

On Ubuntu VMs : `/var/log/syslog`

On CentOS, RHEL and SLES: `/var/log/messages`

When logging to host is enabled, the logs will be stored in `vmware.log` of respective `vmfs` directory of VMs on ESXi host.

`<level>`: **`[1-7]`**, where **4** is for logging level INFO, **7** is for logging level DEBUG.

`<sub-component-name>`: one or more of transport, timer, file, network, process, system

Example:

Enabling the following commands only print logs from that context.

Debug logging for network introspection can be enabled using the following command.

```
/etc/vsep/vsepd refresh-logging 1 7 network
```

Debug logging for process introspection:

```
/etc/vsep/vsepd refresh-logging 1 7 process
```

Debug logging for anti-virus use case:

```
/etc/vsep/vsepd refresh-logging 1 7 file
```

For command processing in timer context (all use cases):

```
/etc/vsep/vsepd refresh-logging 1 7 timer
```

For user monitoring:

```
/etc/vsep/vsepd refresh-logging 1 7 system
```

For framework communication between SVM and Context Mux (all use cases):

```
/etc/vsep/vsepd refresh-logging 1 7 transport
```

Troubleshooting Thin Agent Crashes on Linux

Thin agent dumps core when it crashes. However, it is dependent on the operating system configuration for core dumps. Each Linux distro has different ways and configurations of generating core dump when a system crashes.

For example, you can use `apport` for applications to dump core on a crash, where as in Red Hat you use `abrtd`. However, thin agent dumps backtrace in `/var/log/syslog` (Ubuntu) or `/var/log/messages` (CentOS, RHEL and SLES) in VM or `vmware.log` if logging is enabled on host, depending on the logging destination.

A sample backtrace:

```
localhost systemd: Started Session 4 of user root.
localhost vsep: EMERG: 0: sig_handler(): Received signal: 11
localhost vsep: EMERG: 0: sig_handler(): backtrace returned 7 pointers
localhost vsep: EMERG: 0: sig_handler(): /usr/sbin/vsep(+0x1d35e) [0x7fa2e4c9135e]
localhost vsep: EMERG: 0: sig_handler(): /lib64/libc.so.6(+0x35a00) [0x7fa2e3d76a00]
localhost vsep: EMERG: 0: sig_handler(): /usr/sbin/vsep(+0x3f789) [0x7fa2e4cb3789]
localhost vsep: EMERG: 0: sig_handler(): /lib64/libglib-2.0.so.0(+0x6e0fc) [0x7fa2e47960fc]
localhost vsep: EMERG: 0: sig_handler(): /lib64/libglib-2.0.so.0(+0x6d745) [0x7fa2e4795745]
localhost vsep: EMERG: 0: sig_handler(): /lib64/libpthread.so.0(+0x7df3) [0x7fa2e4109df3]
localhost vsep: EMERG: 0: sig_handler(): /lib64/libc.so.6(clone+0x6d) [0x7fa2e3e373dd]
localhost vsep: EMERG: 0: sig_handler(): Unmarking all fanotify marked mount points
```

NSX Guest Introspection Platform Thin Agent Logs on Windows

The thin agent is installed on the VM Guest OS and detects user activity details.

Log Path and Sample Message

The thin agent consists of GI drivers – `vsepflt.sys`, `vnetwfp.sys` (Windows 10 and later).

The thin agent logs are on the ESXi host, as part of the vCenter Log Bundle. The log path is `/vmfs/volumes/<datastore>/<vmname>/vmware.log` For example: `/vmfs/volumes/5978d759-56c31014-53b6-1866abaace386/Windows10-(64-bit)/vmware.log`

Thin agent messages follow the format of <timestamp> <VM Name><Process Name><[PID]>: <message>.

In the log example below `Guest: vnet or Guest:vsep`, indicate log messages related to the respective GI drivers, followed by debug messages.

For example:

```
2017-10-17T14:25:19.877Z| vcpu-0| I125: Guest: vnet: AUDIT: DriverEntry :
 vnetFilter build-4325502 loaded
2017-10-17T14:25:20.282Z| vcpu-0| I125: Guest: vsep:
AUDIT: VFileSocketMgrConnectHelper : Mux is connected
2017-10-17T14:25:20.375Z| vcpu-0| I125:
Guest: vsep: AUDIT: DriverEntry : vfileFilter build-4286645 loaded

2017-10-17T18:22:35.924Z| vcpu-0| I125: Guest: vsep: AUDIT:
VFileSocketMgrConnectHelper : Mux is connected
```

```
2017-10-17T18:24:05.258Z| vcpu-0| I125: Guest: vsep: AUDIT:
VFileFltPostOpCreate : File (\Windows\System32\Tasks\Microsoft\Windows\
SoftwareProtectionPlatform\SvcRestartTask) in a transaction, ignore
```

Enabling NSX Guest Introspection Platform Thin Agent Driver Logging

Because the debug setting can flood the vmware.log file to the point that it throttles, we recommend you disable the debug mode as soon as you have collected all the required information.

This procedure requires you to modify the Windows registry. Before you modify the registry, ensure to take a backup of the registry. For more information on backing up and restoring the registry, see the Microsoft Knowledge Base article 136393.

1 Click **Start > Run**. Enter regedit, and click **OK**. The Registry Editor window opens. For more information seen the Microsoft Knowledge Base article 256986.

2 Create this key using the registry editor: `HKEY_LOCAL_Machine\SYSTEM\CurrentControlSet\services\vsepflt\parameters`.

3 Under the newly created parameters key, create these DWORDs. Ensure that hexadecimal is selected when putting in these values:

```
Name: log_dest
Type: DWORD
Value: 0x2

Name: log_level
Type: DWORD
Value: 0x10
```

Other values for `log level` parameter key:

```
Audit 0x1
Error 0x2
Warn 0x4
Info 0x8
Debug 0x10
```

4 Open a command prompt as an administrator. Run these commands to unload and reload the NSX Endpoint filesystem mini driver:

- `fltmc unload vsepflt`
- `fltmc load vsepflt`

You can find the log entries in the vmware.log file located in the virtual machine.

Enabling NSX Guest Introspection Platform Driver Logging

Because the debug setting can flood the `vmware.log` file to the point that it can make it to throttle, we recommend you disable the debug mode as soon as you have collected all the required information.

This procedure requires you to modify the Windows registry. Before you modify the registry, ensure to take a backup of the registry. For more information on backing up and restoring the registry, see the Microsoft Knowledge Base article 136393.

1 Click **Start > Run**. Enter regedit, and click **OK**. The Registry Editor window opens. For more information seen the Microsoft Knowledge Base article 256986.

2 Edit the registry:

```
Windows Registry Editor Version 5.0
[HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\vnetwfp\Parameters]
"log_level" = DWORD: 0x0000001F
"log_dest"  = DWORD: 0x00000001
```

3 Reboot the virtual machine.

`vsepflt.sys` Log File Location

With the `log_dest` registry settings **`DWORD: 0x00000001`**, the endpoint thin agent driver logs into the debugger. Run the debugger (DbgView from SysInternals or windbg) to capture the debug output.

Alternatively, you can set the `log_dest` registry setting to **`DWORD:0x000000002`**, in which case the driver logs will be printed to `vmware.log` file, which is located in the corresponding virtual machine folder on the ESXi Host.

Enabling UMC logging

The Endpoint Protection user-mode component (UMC) runs within the VMware Tools service in the protected virtual machine.

1 On Windows XP and Windows Server 2003, create a `tools config` file if it doesn't exist in the following path: `C:\Documents and Settings\All Users\Application Data\VMware\VMware Tools\tools.conf`.

2 On Windows Vista, Windows 7 and Windows Server 2008, create a `tools config` file if it doesn't exist in the following path: `C:\ProgramData\VMWare\VMware Tools\tools.conf`

3 Add these lines in the `tools.conf` file to enable UMC component logging.

```
[logging]
log = true
vsep.level = debug
vsep.handler = vmx
```

With the `vsep.handler = vmx` setting, the UMC component logs into the `vmware.log` file, which is located in the corresponding virtual machine folder on the ESXi host.

With the following setting logs, the UMC component logs will be printed in the specified log file.

```
vsep.handler = file
vsep.data = c:/path/to/vsep.log
```

Troubleshooting the Thin Agent on Windows

1 Check the compatibility of all the components involved. You need the build numbers for ESXi, vCenter Server, NSX Manager, and the Security solution you have selected (for example, Trend Micro, McAfee, Kaspersky, or Symantec). After this data is collected, you can compare the compatibility of the vSphere components. For more information, see the VMware Product Interoperability Matrices.

2 Ensure that VMware Tools™ is up-to-date. If you see that only a particular virtual machine is affected, see Installing and upgrading VMware Tools in vSphere (2004754).

3 Verify that the thin agent is loaded by running the PowerShell command `fltmc`.

Verify that vsepflt is included in the list of drivers. If the driver is not loaded, try loading the driver with the `fltmc load vsepflt` command.

4 If the thin agent is causing a performance problem with the system, unload the driver with this command: `fltmc unload vsepflt`.

5 Next, perform a test to get a baseline. You can then load the driver and perform another test by running this command:

`fltmc load vsepflt`.

If you do verify that there is a performance problem with the Thin agent, see Slow VMs after upgrading VMware tools in NSX and vCloud Networking and Security (2144236).

6 If you are not using Network Introspection, remove or disable this driver.

Network Introspection can also be removed through the Modify VMware Tools installer:

a Mount the VMware Tools installer.

b Navigate to **Control Panel > Programs and Features**.

c Right-click **VMware Tools > Modify**.

d Select **Complete install**.

e Find NSX File Introspection. This contains a subfolder for Network Introspection.

f Disable **Network Introspection**.

g Reboot the VM to finish the uninstallation of the driver.

7 Enable debug logging for the thin agent. All debugging information is configured to log to the vmware.log file for that virtual machine.

8 Review the file scans of the thin agent by reviewing the procmon logs. For more information, see Troubleshooting vShield Endpoint performance issues with anti-virus software (2094239).

Troubleshooting Thin Agent Crashes on Windows

If the Thin Agent crashes, the core file is generated in the `/directory`. Collect the core dump file (core) from `location / directory`.

Collect Environment and Workload Details

Collect ESXi version, OS version and logs generated by the various components of the NSX Guest Introspection Platform.

Collect Environment and Workload Details

1 Determine if Guest Introspection is used in your environment. If it is not, remove the Guest Introspection service for the virtual machine, and confirm that the problem is resolved. Troubleshoot a Guest Introspection problem only if Guest Inspection is required.

2 Collect environment details:

   a To collect the ESXi build version, run the command `uname –a` on the ESXi host or select a host in the vSphere Web Client and look for the build number at the top of the right pane.

   b Linux or Windows product version and build number.

   c */usr/sbin/vsep -v* returns the production version:

```
Build number
------------------
Ubuntu
dpkg -l | grep vmware-nsx-gi-file
SLES12 and RHEL7
rpm -qa | grep vmware-nsx-gi-file
```

3 Collect the NSX for vSphere version, and the following:

- Partner solution name and version number
- EPSec Library version number used by the partner solution: Log into the SVM and run `strings <path to EPSec library>/libEPSec.so | grep BUILD`
- Guest operating system in the virtual machine
- Any other third-party applications or file system drivers

4 ESX GI Module (MUX) version - run the command *esxcli software vib list | grep nsx-context-mux.*

5 Collect workload details, such as the type of server.

6 Collect ESXi host logs. For more information, see Collecting diagnostic information for VMware ESX/ESXi (653).

7 Collect logs from the consumers of NSX Guest Introspection Platform. The consumers are Endpoint Protection (service virtual machine), NSX Malware Prevention (Security Hub VM), NSX Intelligence and IDFW. Contact your partner for more details on SVM log collection.

8 Collect the VMware `vmss` file of the virtual machine in a suspended state, see Suspending a virtual machine on ESX/ESXi to collect diagnostic information (2005831), or crash the virtual machine and collect the full memory dump file. VMware offers a utility to convert an ESXi `vmss` file to a core dump file. See Vmss2core fling for more information.

Supported Software

Guest Introspection is interoperable with specific versions of software.

VMware Tools

Check out interoperability between VMware Tools and NSX. See VMware Product Interoperability Matrices. For the list of supported Windows operating system versions, see the VMware Tools release notes at https://docs.vmware.com/en/VMware-Tools/index.html.

Supported Hosts

For supported ESXi hosts, see the VMware Product Interoperability Matrices.

Security Monitoring

Using Aria Operations for Logs for Unified Security Logs

With Aria Operations for Logs you can view the security flow logs of the NSX environment.

The following security features support flow logging:

- TLS Inspection
- Gateway IDPS
- URL Filtering

Unified Security Logs

All the security verticals generate and save unified security flow logs in the Unified Security Logs format in a single log file on a node. This single log is exported to syslog server, which is configured for Aria Operations for Logs. Aria Operations for Logs will then process the logs to provide further log management, analysis, and display them by using NSX Content Pack.

Displaying Logs on Aria Operations for Logs

The dashboard 'NSX - Unified Security Flow Logs' in the NSX Content Pack shows chart widgets, which are visual representation of the security flow logs.

An Aria Operations for Logs Content Pack is a plugin. It contains dashboards, extracted fields, saved queries, and alerts that are related to a specific product or set of logs.

The NSX Content Pack is available on Aria Operations for Logs Marketplace.

For more information about Aria Operations for Logs and how to install a Content Pack from Content Pack Market place, see the chapter *Install a Content Pack from the Content Pack Marketplace* from the Aria Operations for Logs documentation.

Top-N and Last X-hours

You can also query events in Aria Operations for Logs for Top-N information from last X-hours using Interactive Analytics and Content Pack.

Remote Logging Server

To send logs to a remote logging server, NSX appliances and hypervisors must be configured with remote logging on each node separately.

Note For the logs to be sent to the syslog server, you must enable Logging for the specific Rules on NSX Manager.

For more information, see Configure Remote Logging.

If logs are not received by the remote log server, see Troubleshooting Syslog Issues.

Unified Security Log Format

You can also get the flow logs at `/var/log/syslog | grep 'unified-logs'` on the edges. For example, for

TLS Inspection:


```
2021-10-12T09:00:46.192Z nsxedge-18734920-1-mps29 NSX 22621 SYSTEM [nsx@6876 comp="nsx-edge"
subcomp="tls-proxy" s2comp="unified-logs" level="INFO"] {"event_type": "fw-flow-terminate-
log",
"event_trigger": ["fw-rule-log"], "origin": {"fw_type": "gateway", "fw_uuid":
"79427614-4a0d-2692-032c-eb4692f717a9",
"node_uuid": "ec11a626-f425-3bc2-671d-a656500003b2"}, "flow": {"start":
"2021-10-12T09:00:46.723Z",
"end": "2021-10-12T09:00:46.773Z", "ip_ver": "ipv4", "flow_id": "0x1e0000704f000018",
"src_ip": "192.168.100.160", "src_port": 25700, "dest_ip": "1.1.5.10", "dest_port": 443,
"proto": "TCP", "tcp_flags": "",
"bytes_toserver": 95, "bytes_toclient": 29873, "reason": "FIN-close", "final_action":
"PASS"}, "fw": {"action": "PASS", "rule_id": 1002,
"direction": "", "rule_tag": ""}, "http": {"http_method": "", "hostname": "www.facebook.com",
"url": "www.facebook.com/benign_pdf1.pdf", "scheme": "", "http_user_agent": "", "status": "",
"site_category": ""SOCIAL_NETWORK"", "site_reputation": "Trustworthy"},"tls_inspection":
{"action": "PASS", "rule_id": 1008, "domain": "www.facebook.com", "cert_status": "ok",
"tls_version_toserver": "TLSv1.2",
"cipher_to_server": "TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256", "reason": "", "tls_rule_tag":
"TLS External Rule"}, "idps": {"action": "PASS", "rule_id": 1007,
"ids_profile_id": "00000000-0000-0000-0000-000000000000", "alert_event": ["SOCIAL_NETWORK"],
"protocol_event": ["http"]}, "app_id": {"app": ""APP_HTTP", "APP_FACEBOOK", "APP_SSL""}
```

Gateway IDPS:


```
2021-11-11T04:07:37.489Z nsxedge-18866667-2-NAT-fc-3-nov NSX 27330 SYSTEM [nsx@6876 comp="nsx-
edge" subcomp="datapathd" s2comp="unified-logs" level="INFO"]
 {"event_type": "fw-flow-terminate-log", "event_trigger": ["ids-rule-log"], "origin":
{"fw_type": "gateway", "fw_uuid": "544ee558-fad0-e624-019f-e8e3e0472c91",
"node_uuid": "dabf16c9-ec11-833c-0c00-8c832f771729"}, "flow": {"start":
"2021-11-11T04:07:07.000Z", "end": "2021-11-11T04:07:37.000Z",
"ip_ver": "ipv4", "flow_id": "0xd44d00306e0a0004", "src_ip": "1.1.1.10", "src_port": 35714,
"dest_ip": "10.142.7.1", "dest_port": 53,
"proto": "UDP", "tcp_flags": "FPW", "bytes_toserver": 297878, "bytes_toclient": 71,
"pkts_toserver": 71, "pkts_toclient": 1,
"reason": "FIN-close", "final_action": "PASS"}, "fw": {"action": "PASS", "rule_id": 2025,
"direction": "", "rule_tag": ""},
"l7profile": {"entry_id": "00000000-0000-0000-0000-000000000000", "action": "PASS"},
"http": {"http_method": "", "hostname": "", "url": "", "scheme": "", "http_user_agent": "",
"status": "", "site_category": "",
"site_reputation": "UNKNOWN"}, "idps": {"action": "IDP_DETECT", "rule_id": 2028,
"ids_profile_id": "9872e27a-ead4-4c93-af5f-4df1ec0c73e1",
"alert_event": [], "protocol_event": []}, "app_id": {"app": ""APP_DNS""}}
```


URL Filtering:


```
2021-11-08T08:30:59.208Z nsxedge-18866667-2-NAT-fc-3-nov NSX 9495 SYSTEM [nsx@6876 comp="nsx-
edge" subcomp="datapathd" s2comp="unified-logs" level="INFO"]
{"event_type": "fw-flow-terminate-log", "event_trigger": ["fw-rule-log"], "origin":
{"fw_type": "gateway", "fw_uuid": "544ee558-fad0-e624-019f-e8e3e0472c91",
"node_uuid": "dabf16c9-ec11-833c-0c00-8c832f771729"}, "flow": {"start":
"2021-11-08T08:30:57.000Z", "end": "2021-11-08T08:30:59.000Z",
"ip_ver": "ipv4", "flow_id": "0x4001006c04000000", "src_ip": "1.1.1.10", "src_port": 43600,
"dest_ip": "13.226.234.18",
"dest_port": 80, "proto": "TCP", "tcp_flags": "FREW", "bytes_toserver": 54968,
"bytes_toclient": 444,
"pkts_toserver": 444, "pkts_toclient": 7, "reason": "FIN-close", "final_action": "PASS"},
"fw": {"action": "PASS", "rule_id": 1004, "direction": "",
"rule_tag": ""}, "l7profile": {"entry_id": "e9580107-2749-471c-be82-715d530bf4d4", "action":
"PASS"},
"http": {"http_method": "", "hostname": "", "url": "espn.com/", "scheme": "",
"http_user_agent": "", "status": "",
"site_category": ""SPORTS"", "site_reputation": "TRUSTWORTHY"}, "app_id": {"app":
""APP_HTTP", "APP_ESPN""}}
```


Monitoring Security Statistics

NSX Application Platform collects and stores statistics for security features. You can view these metrics by invoking the time series metrics APIs.

Before you begin

You must have the "NSX Gateway Firewall with Advanced Threat Prevention" license for time series monitoring.

You must deploy NSX Application Platform. For more details about deploying NSX Application Platform, see the *Deploying and Managing NSX Application Platform* guide. The Metrics feature is enabled by default when you deploy NSX Application Platform.

Security Statistics

The following security features generate statistics with API/CLI respectively:

- TLS Inspection
- Gateway IDPS
- Gateway Firewall and Connections

Time series metrics is available only for TLS Inspection, Gateway IDPS, and Gateway Firewall. You can retrieve these metrics through the metric APIs.

Note For URL Filtering and Malware prevention, only point-in-time security metrics are available and displayed on the NSX Manager user interface.

Metrics API

You can use Metrics APIs to fetch the time series metrics. These APIs can take multiple intent paths or UUIDS as input for a specific resource type, such as edge or firewall, and return the corresponding metrics.

Using time series metrics, you can monitor the trend in key performance indicators, detect anomalies, perform before and after analysis, and get the historical context which can help in troubleshooting.

Based on your role, you can view metrics of only those objects for which you have the authorization.

For high-level information about the time series metrics workflow, see APIs to Fetch Time-Series Metrics. For complete information about how to invoke the time series metrics APIs, see *NSX Intelligence & NSX Application Platform API Guide*.

Display Statistics on NSX Manager User Interface.

Two types of metrics can be viewed on the NSX Manager user interface:

- Point-in-time - Recent data fetched from the Node.
- Time series - Historical data to provide daily, weekly, monthly and yearly view.

All time series metrics are accessible through the metrics API. However, some of the time series metrics are also available on the NSX Manager user interface.

For dashboard information associated to each security feature, see Security Overview.

For information on how to monitor NSX Edge Nodes and Gateways, see Monitor NSX Edge Nodes and Gateways.

Security Terminology

The following terms are used throughout distributed firewall.

Table 16-1. Security-Related Terminology

Construct	Definition
Policy	A security policy includes various security elements including firewall rules and service configurations. Policy was previously called a firewall section.
Rule	A set of parameters with which flows are evaluated against, and define which actions will be taken upon a match. Rules include parameters such as source and destination, service, context profile , logging, and tags.
Group	Groups include different objects that are added both statically and dynamically, and can be used as the source and destination field of a firewall rule. Groups can be configured to contain a combination of virtual machines, IP sets, MAC sets, logical ports, logical switches, AD user groups, and other nested groups. Dynamic inclusion of groups can be based on tag, machine name, OS name, or computer name. When you create a group, you must include a domain that it belongs to, and by default this is the default domain. Groups were previously called NSGroup or security group.
Service	Defines a combination or port and protocol. Used to classify traffic based on port and protocol. Pre-defined services and user-defined services can be used in firewall rules.
Context Profile	Defines context aware attributes including APP-ID and domain name. Also includes sub attributes such as application version, or cipher set. Firewall rules can include a context profile to enable Layer-7 firewall rules.

Identity Firewall

With Identity Firewall (IDFW) features an NSX administrator can create Active Directory user-based distributed firewall (DFW) rules.

IDFW can be used for Virtual Desktops (VDI), Remote desktop sessions (RDSH support), and physical machines, enabling simultaneous log ins by multiple users, user application access based on requirements, and the ability to maintain independent user environments. VDI management systems control what users are granted access to the VDI virtual machines. NSX controls access to the destination servers from the source virtual machine (VM), which has IDFW enabled. With RDSH, administrators create security groups with different users in Active Directory (AD), and allow or deny those users access to an application server based on their role. For example, Human Resources and Engineering can connect to the same RDSH server, and have access to different applications from that server.

IDFW must know which desktop an Active Directory (AD) user logs onto in order to apply firewall rules. There are two methods IDFW uses for logon detection: Guest Introspection (GI) and/or event log scraping. Guest Introspection is deployed on ESXi clusters where IDFW virtual machines are running. When network events are generated by a user, a guest agent installed on the VM forwards the information through the Guest Introspection framework to the NSX

Manager. The second option is the Active Directory event log scraper. Event log scraping enables IDFW for physical devices. Configure the Active Directory event log scraper in the NSX Manager to point at an instance of your Active Directory domain controller. NSX Manager will then pull events from the AD security event log.

Event log scraping can be used for virtual machines, however when both the AD log scraper and Guest Introspection are used, Guest Introspection will take precedence over event log scraping. Guest introspection is enabled through VMware Tools, and if you are using the complete VMware Tools installation and IDFW, guest introspection will take precedence over event log scraping.

IDFW can also be used on VMs that have supported operating systems. See Identity Firewall Supported Configurations.

IDFW processes the user identity at the source only in firewall rules. Identity-based groups cannot be used as the destination in distributed firewall and gateway firewall rules.

Note IDFW relies on the security and integrity of the guest operating system. There are multiple methods for a malicious local administrator to spoof their identity to bypass firewall rules. User identity information is provided by the NSX Guest Introspection Thin Agent inside guest VMs. Security administrators must ensure that thin agent is installed and running in each guest VM. Logged-in users should not have the privilege to remove or stop the agent.

For supported IDFW configurations see Identity Firewall Supported Configurations.

Note that IDFW rules are not supported on Global Managers in a Federation environment. IDFW can still be used locally in Federated sites by creating IDFW rules on Local Managers.

Identity Firewall Workflow

IDFW enhances traditional firewall by allowing firewall rules based on user identity. For example, administrators can allow or disallow customer support staff to access an HR database with a single firewall policy.

Identity based firewall rules are determined by membership in an Active Directory (AD) group membership. See Identity Firewall Supported Configurations.

IDFW processes the user identity at the source only in firewall rules. Identity-based groups cannot be used as the destination in distributed firewall and gateway firewall rules.

Note For Identity Firewall rule enforcement, Windows Time service should be **on** for all VMs using Active Directory. This ensures that the date and time is synchronized between Active Directory and VMs. AD group membership changes, including enabling and deleting users, do not immediately take effect for logged in users. For changes to take effect, users must log out and then log back in. AD administrator's should force a logout when group membership is modified. This behavior is a limitation of Active Directory.

Prerequisites

If Windows auto-logon is enabled on VMs, go to **Local Computer Policy > Computer configuration > Administrative Templates > System > Logon** and enable **Always wait for the network at computer startup and logon**.

For supported IDFW configurations see Identity Firewall Supported Configurations.

Procedure

1. Enable NSX File Introspection driver and NSX Network Introspection driver (VMware Tools full installation adds these by default), or event log scraping. See Identity Firewall Event Log Sources.

   Event log scraping enables IDFW for physical devices. Event log scraping can be used for virtual machines, however guest introspection will take precedence over event log scraping. Guest Introspection is enabled through VMware Tools and if you are using the complete VMware Tools installation and IDFW, guest introspection will take precedence over event log scraping.
2. Enable Identity Firewall on DFW and GFW.
3. Configure Active Directory (required) and event log scraping (optional) Configuring Active Directory and Event Log Scraping.
4. Configure Active Directory sync operations: Synchronize Active Directory.
5. Create a group with Active Directory group members: Add a Group.
6. Assign group with AD group members to a distributed firewall rule or gateway firewall rule. If creating a DFW rule using guest introspection, make sure that the **Applied to** field applies to the destination group: Add a Distributed Firewall . The **Source** field should be an AD based group.

Enable Identity Firewall on DFW and GFW

Identity Firewall must be activated for IDFW firewall rules to take effect.

Procedure

1. Select **Security > Distributed Firewall** or **Security > Gateway Firewall**
2. In the right corner, click **Actions > General Setting**.
3. Toggle the status button to activate IDFW.

   Distributed firewall or gateway firewall must also be activated for IDFW to work.
4. For distributed firewall, to enable IDFW on standalone hosts or clusters, select the **Identity Firewall Settings** tab. If you are providing IDFW for physical machines you also need to turn on Event Log sources, see Identity Firewall Event Log Sources.
5. Toggle the **Enable** bar, and select the standalone hosts, or select the cluster where the IDFW host must be activated.

6 Click **Save**.

Identity Firewall Best Practices

The following best practices will help maximize the success of identity firewall rules.

- IDFW supports the following protocols:
  - Single user (VDI, or Non-RDSH Server) use case support - TCP, UDP, ICMP
  - Multi-User (RDSH) use case support - TCP, UDP

  .
- Any change on a domain, including a domain name change, will trigger a full sync with Active Directory. Because a full sync can take a long time, we recommend syncing during off-peak or non-business hours.
- For local domain controllers, the default LDAP port 389 and LDAPS port 636 are used for the Active Directory sync, and should not be edited from the default values.

Identity Firewall Supported Configurations

The following configurations are supported for IDFW on virtual machines (VMs).

IDFW supports the following protocols:

- Single user (VDI, or Non-RDSH Server) use case support - TCP, UDP

  Note ICMP filtering can be enabled with VMware Tools 12.x. For more information, see KB articles 79185 and 88273.

- Multi-User (RDSH) use case support - TCP, UDP

Guest Operating Systems	Enforcement Type
Windows 8	Desktop - supports desktop users use case
Windows 10	Desktop - supports desktop users use case
Windows 2012	Server - supports server users use case
Windows 2012R2	Server - supports server users use case
Windows 2016	Server - supports server users use case
Windows 2019	Server - supports server users use case
Windows 2022	Server - supports server users use case
Windows 2012R2	RDSH - supports Remote Desktop Session Host
Windows 2016	RDSH - supports Remote Desktop Session Host
Windows 2019	RDSH - supports Remote Desktop Session Host
Windows 2022	RDSH - supports Remote Desktop Session Host

Active Directory Domain Controllers:

- Windows Server 2012

- Windows Server 2012R2
- Windows Server 2016
- Windows Server 2019
- Windows Server 2022

Host operating system: ESXi

VMware Tools - Version 11

- VMCI Driver
- NSX File Introspection Driver
- NSX Network Introspection Driver

Layer 7 Context Profile

Layer 7 App IDs are configured as part of a context profile.

A context profile can specify one or more App IDs, and can also include sub-attributes, for use in distributed firewall (DFW) rules and gateway firewall rules. When a sub-attribute, such as TLS version 1.2 is defined, multiple application identity attributes are not supported. In addition to attributes, DFW also supports a Fully Qualified Domain Name (FQDN) or URL that can be specified in a context profile for FQDN allowlisting or denylisting. See FQDN Filtering for more information. FQDNs can be configured with an attribute in a context profile, or each can be set in different context profiles. After a context profile has been defined, it can be applied to one or more distributed firewall rules.

Note

- Gateway firewall rules do not support the use of FQDN attributes or other sub attributes in context profiles.
- Context profiles are not supported on tier-0 gateway firewall policy.

When a context profile is used in a rule, any traffic coming in from a virtual machine is matched against the rule table based on the 5-tuple. If the rule matching the flow also includes a Layer 7 context profile, the packet is redirected to a user-space component called the vDPI engine. Subsequent packets are sent to the vDPI engine for each flow. After the App ID has been determined, the information is stored in the in-kernel context table. When the next packet for the flow comes in the information in the context table is compared with the rule table again, and is matched on 5-tuple and the Layer 7 App ID. The appropriate action as defined in the fully matched rule is taken, and if there is an ALLOW rule, all subsequent packets for the flow are processed in the kernel, and matched against the connection table. For fully matched DROP rules, a reject packet is generated. Logs generated by the firewall will include the Layer 7 App ID and applicable URL, if that flow was sent to the vDPI engine.

Rule processing for an incoming packet:

1 Upon entering a DFW or Gateway filter, packets are looked up in the flow table based on 5-tuple.

2 If no flow/state is found, the flow is matched against the rule-table based on 5-tuple and an entry is created in the flow table.

3 If the flow matches a rule with a Layer 7 service object, the flow table state is marked as “DPI In Progress.”

4 The traffic is then punted to the DPI engine. The DPI Engine determines the App ID.

5 After the App ID has been determined, the DPI Engine sends down the attribute which is inserted into the context table for this flow. The "DPI In Progress" flag is removed, and traffic is no longer punted to the DPI engine.

6 The flow (now with App ID) is reevaluated against all rules that match the App ID, starting with the original rule that was matched based on 5-tuple, and the first fully matched L4/L7 rule is picked up. The appropriate action is taken (allow/deny/reject) and the flow table entry is updated accordingly.

Layer 7 Firewall Rule Workflow

Layer 7 App Ids are used in creating context profiles, which are used in distributed firewall rules or gateway firewall rules. Rule enforcement based on attributes enables users to allow or deny applications to run on any port.

NSX provides built in App IDs for common infrastructure and enterprise applications. App Ids include versions (SSL/TLS and CIFS/SMB) and Cipher Suite (SSL/TLS). For distributed firewall, App Ids are used in rules through context profiles, and can be combined with FQDN allowlisting and denylisting.

Note

- Gateway firewall rules do not support the use of FQDN attributes or other sub attributes in context profiles.
- Context profiles are not supported on tier-0 gateway firewall policy.

Supported App IDs and FQDNs:

- For FQDN, users need to configure a high priority rule with a DNS App ID for the specified DNS servers on port 53.
- SYSLOG App ID is detected only on standard ports.

Note that if you are using a combination of Layer 7 and ICMP, or any other protocols you need to put the Layer 7 firewall rules last. Any rules after a Layer 7 any/any rule will not be executed.

Procedure

1 Create a custom context profile: Profiles.

2 Use the context profile in a distributed firewall rule, or a gateway firewall rule: Add a Distributed Firewall or Add a Gateway Firewall Policy and Rule.

Multiple App ID context profiles can be used in a firewall rule with services set to **Any**. For ALG profiles (FTP, ORACLE, DCERPC, TFTP), one context profile is supported per rule.

Distributed Firewall

Distributed firewall comes with predefined categories for firewall rules. Categories allow you to organize security policies.

Categories are evaluated from left to right (Ethernet > Emergency > Infrastructure > Environment > Application), and the distributed firewall rules within the category are evaluated top down.

Table 16-2. Distributed Firewall Rule Categories

Ethernet	Emergency	Infrastructure	Environment	Application
We recommend you include Layer 2 rules for this category.	We recommend you include quarantine and allow rules for this category.	We recommend you include rules which define access to shared services for this category. For example: ■ AD ■ DNS ■ NTP ■ DHCP ■ Backup ■ Management servers	We recommend you include rules between zones for this category. For example: ■ Production vs development ■ PCI vs non-PCI ■ Inter business unit rules	We recommend you include rules between: ■ Applications ■ Application tiers ■ Micro services

FQDN Filtering

Set up a distributed firewall rule to filter specific domains identified with a fully qualified domain name, for example, **.office365.com*.

You must set up a DNS rule first, and then the FQDN allowlist or denylist rule below it. NSX uses time to live (TTL) in the DNS response (coming from DNS server to the virtual machine), for keeping the DNS to IP mapping cache entry for the virtual machine. To override the DNS TTL using a DNS security profile, see Configure DNS Security. For FQDN filtering to be effective, virtual machines need to use a DNS server for domain resolution (no static DNS entries), and also need to honor the TTL received in the DNS response. NSX uses DNS Snooping to obtain a mapping between the IP address and the FQDN.

This feature works at layer 7 and does not cover ICMP. If a user creates a denylist rule for all services on `example.com` the feature is working as intended if ping `example.com` responds, but curl `example.com` does not.

Selecting a wild card FQDN is a best practice because it includes sub domains. For example, selecting `*example.com`, would include sub domains such as `americas.example.com` and `emea.example.com`. Using `example.com` would not include any sub domains.

FQDN-based rules are retained during vMotion for ESXi hosts.

Note FQDN filtering is available only with TCP and UDP traffic.

Prerequisites

To use a user-defined FQDN, see FQDNs.

Create a DNS rule if it doesn't already exist:

1 Navigate to **Security > Distributed Firewall**.

2 Select the check box next to a policy section and click **Add Rule**.

3 Provide a name for the firewall rule, such as **`DNS rule`**, and provide the following details:

Variable	Description
Name	Provide a name for the rule, such as L7 DNS Rule
Source	Any or specific group
Destination	Any or specific group
Services	Click the edit icon and select the **DNS-TCP** and **DNS-UDP** service as applicable to your environment.
Context Profiles	Click the edit icon and select the **`DNS`** context profile. This is system generated context profile, and is available in your deployment by default.
Applied To	Select a group as required.
Action	Select **Allow**.

4 Click **Publish**.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Navigate to **Security > Distributed Firewall**.

3 Click **Add Rule** to set up the FQDN allowlisting or denylisting rule.

4 Name the rule appropriately, such as, **`FQDN/URL Allowlist`**.

5 Provide the following details:

Option	Description
Services	Click the edit icon and select the service you want to associate with this rule, for example, HTTP.
Context Profiles	Click the edit icon, and **Add Context Profile** and name the profile. In the Attributes column, select **Set > Add Attribute > Domain (FQDN) Name** . Select the list of Attribute Name/Values from the predefined list, or create a custom FQDN. See Context Profilesfor details. Click **Add**, and **Apply**.
Applied To	Select DFW or a group as required.
Action	Select **Allow**, **Drop**, or **Reject**.

6 Click **Publish**.

Firewall Drafts

A draft is a complete distributed firewall configuration with policy sections and rules. Drafts can be auto saved or manually saved, and immediately published or saved for publishing at a later date.

To save a manual draft firewall configuration, go to the upper right of the distributed firewall screen and click **Actions > Save**. After saving, the configuration can be viewed by selecting **Actions > View**. Auto drafts are enabled by default. Auto drafts can be disabled by going to **Actions > General Settings**. When auto drafts are enabled, any changes to a firewall configuration results in a system generated autodraft. A maximum of 100 auto drafts and 10 manual drafts can be saved. Auto drafts can be edited and saved as a manual draft, for publishing now or later. To prevent multiple users from opening and editing the draft, manual drafts can be locked. When a draft is published, the current configuration is replaced by the configuration in the draft.

Save or View a Firewall Draft

A draft is a distributed firewall configuration that has been published, or saved for publishing at a later date. Drafts are created automatically, and manually.

Manual drafts can be edited and saved. Auto drafts can be cloned, and saved as manual drafts, and then edited. The maximum number of drafts that can be saved is 100 autodrafts and 10 manual drafts.

Procedure

1 Click **Security > Distributed Firewall**.

2 To save a firewall configuration manually, go to **Actions > Save**.

A manual draft can be saved, or edited and then saved. After saving, you can revert to the original configuration.

3 **Name** the configuration.

4 To prevent multiple users from opening and editing a manual draft, **Lock** the configuration, and add a comment.

5 Click **Save**.

6 To view the saved configuration, click **Actions > View**.

A timeline opens up showing all saved configurations. To see details such as draft name, date, time and who saved it, point to the dot or star icon of any draft. Saved configurations can be filtered by time, showing all drafts in the last one day, one week, 30 days, or the last three months. They can be filtered by aurodraft and saved by me. They can also be filtered by name, by using the search tool on the top right.

7 Hover over a draft to view name, date and time details of the saved configuration. Click the name to view draft details.

The detailed draft view shows the required changes to be made to the current firewall configuration, in order to be in sync with this draft. If this draft is published, all of the changes visible in this view will be applied to the current configuration.

Clicking the downward arrow expands each section, and displays the added, modified, and deleted changes in each section. The comparison shows added rules with a green bar on the left side of the box, modified elements (such as a name change) have a yellow bar, and deleted elements have a red bar.

8 To edit the name or description of a selected draft, click the menu icon (three dots) from the **View Draft Details** window, and select **Edit**.

Manual drafts can be locked. If locked, a comment for the draft must be provided.

Some roles, such as enterprise administrator have full access credentials, and cannot be locked out. See Role-Based Access Control.

9 Auto drafts and manual drafts can also be cloned and saved by clicking **Clone**.

In the Saved Configurations window, you can accept the default name, or edit it. You can also lock the configuration. If locked, a comment for the draft must be provided.

10 To save the cloned version of the draft configuration, click **Save**. The draft is now present in the Saved Configurations section.

What to do next

After viewing a draft, you can load and publish it. It is then the active firewall configuration.

Publish or Revert a Firewall Draft

Both auto drafts and saved manual drafts can be loaded and published to become the active configuration.

During publishing, a new auto draft is created. This auto draft can be published to revert to the previous configuration.

Procedure

1 To view the saved configuration, click **Actions > View**.

   A timeline opens up showing all saved configurations. To see details such as draft name, date, time and who saved it, point to the dot icon of any draft. Saved configurations are filtered by time, showing all drafts created in 1 day, 1 week, 30 days, or the last 3 months.

2 Click a draft name and the View Draft Details window appears.

3 Click **Load**. The new firewall configuration appears on the main window.

   Note A draft cannot be loaded if firewall filters are being used, or if there are unsaved changes in the current configuration.

4 To commit the draft configuration and make it active, click **Publish**. To return to the previous published configuration, click **Revert**.

   After publishing, the changes in the draft will be present in the active configuration.

5 To edit the contents of the selected draft before publishing, after clicking **Load**, edit the configuration.

6 To save the edited version of the draft configuration, click **Actions > Save**.

   Manual drafts can be saved as a new configuration, or an update to the existing configuration. Auto drafts can only be saved as a new configuration.

7 Enter a **Name** , and optional **Description**. You can also **Lock** the draft. If locked, a comment for the draft must be provided.

8 Click **Save**.

9 To commit the draft configuration and make it active, click **Publish**, or to return to the previous published configuration, click **Revert**.

Configure Malicious IP Feeds

For Distributed Firewall, you can setup Malicious IP Feeds and download a list of known malicious IPs.

The system downloads these IPs from NTICS cloud service and creates a malicious IP group with them. You can also create custom malicious IP groups to specify IPs and IP addresses only groups that should be treated as exceptions and must not be blocked. To block access to malicious IPs, configure firewall rules containing malicious IP groups. You can also monitor the system for any exceptions and if required exclude IPs from getting blocked.

Once you activate Malicious IP Feeds, the IPs are updated at a system defined frequency.

Note If you are the Greenfield customer, this feature is by default enabled for you with the appropriate license. If you are the Brownfield customer, you will have to perform the steps mentioned in the procedure to enable this feature.

You can also manually update the IPs by clicking **Download Latest Feed** on the Settings page. Later, at any time if you turn off Malicious IP Feeds and you have rules with malicious IP groups, the rules might get enforced with outdated data.

To activate Malicious IP Feeds:

Procedure

1 Navigate to **Security > Distributed Firewall**.

2 Go to **Actions > General Settings > Malicious IP Feeds**.

3 Set the **Auto Update Malicious IP** toggle to **On**. The **Last Updated** field shows the status of the download. It also shows the date and time of the the last download.

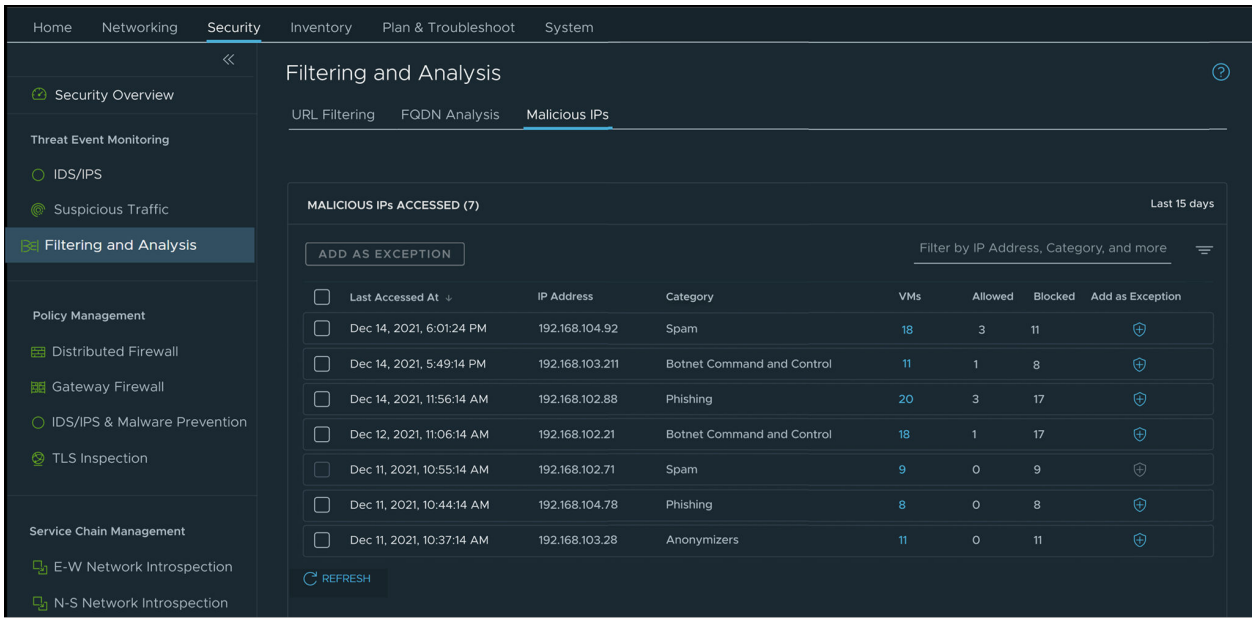
Results

The system downloads malicious IPs and creates a malicious IP group with the downloaded IPs.

Malicious IPs Filtering and Analysis Dashboard

The Filtering and Analysis Dashboard displays all malicious IPs that accessed the system or are accessed by the system. You can also apply filters to view the drilled-down data.

The page displays detailed information about each malicious IPs, such as category that it blocked, total count of VMs that accessed it or accessed by it, number of times the malicious IP is blocked or allowed.



To view details about number of times an IP is blocked or allowed by VMs, click the VM count.

You can also add IPs to the exception list of any malicious IP group by selecting the required IPs and clicking **Add as Exception**.

Add a Distributed Firewall

Distributed firewall monitors all the East-West traffic on your virtual machines.

The procedure in this topic explains the workflow for adding firewall policies that are applied to the NSX Distributed Firewall or to specific groups with NSX-managed objects.

If your NSX environment has Antrea containers registered to it, you can create Distributed Firewall policies and apply them to Antrea container clusters. For more information, see:

- Distributed Firewall Policies for Securing Traffic Within an Antrea Kubernetes Cluster
- Add a Distributed Firewall Policy to Secure Traffic Within an Antrea Kubernetes Cluster

Note NSX does not support mixing the rules created with NSX-managed objects and with Antrea container cluster objects in the same Distributed Firewall policy. In other words, the firewall rules that you apply to NSX Distributed Firewall and to Antrea container clusters must be in separate policies.

Prerequisites

VMs to be DFW-protected must have their vNIC connected to an NSX overlay or VLAN segment. In NSX, distributed firewall protects workloads that are natively connected to a VDS distributed port-group (DVPG). For more information see Distributed Security for vSphere Distributed Switch.

If you are creating rules for Identity Firewall, first create a group with Active Directory members. To view supported protocols for IDFW, see Identity Firewall Supported Configurations. When creating a DFW rule using guest instrospection, make sure that the **Applied to** field applies to the destination group.

Note For Identity Firewall rule enforcement, Windows Time service should be **on** for all VMs using Active Directory. This ensures that the date and time is synchronized between Active Directory and VMs. AD group membership changes, including enabling and deleting users, do not immediately take effect for logged in users. For changes to take effect, users must log out and then log back in. AD administrator's should force a logout when group membership is modified. This behavior is a limitation of Active Directory.

Note that if you are using a combination of Layer 7 and ICMP, or any other protocols you need to put the Layer 7 firewall rules last. Any rules after a Layer 7 any/any rule will not be executed.

For Federation-specific details on distributed firewall policy and rule creation, see Create DFW Policies and Rules from Global Manager.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Security > Distributed Firewall** from the navigation panel.

3 Ensure that you are in the correct pre-defined category, and click **Add Policy**.

For more about categories, see Distributed Firewall .

4 Enter a **Name** for the new policy section.

5 (Optional) Use **Applied to** to apply the rules within policy to a selected group. By default, the policy **Applied to** field is set to DFW, and the policy rules are applied to all workloads. When you change the default, if both the policy level and the rules within have **Applied to** set to a group, the policy level **Applied to** takes precedence over **Applied to** at the rule level.

Note Groups consisting of only IP addresses, MAC Addresses, or Active Directory groups cannot be used in the **Applied To** text box.

Applied to defines the scope of enforcement per policy, and is used mainly for optimization of resources on ESXi host. It helps in defining a targeted policy for specific zones, tenants or applications, without interfering with other policy defined for other applications, tenants and zones.

6 (Optional) To configure the following policy settings, click the gear icon.

Option	Description
TCP Strict	A TCP connection begins with a three-way handshake (SYN, SYN-ACK, ACK) and typically ends with a two-way exchange (FIN, ACK). In certain circumstances, the distributed firewall (DFW) might not see the three-way handshake for a particular flow (due to asymmetric traffic or the distributed firewall being enabled while a flow exists). By default, the distributed firewall does not enforce the need to see a three-way handshake, and picks up sessions that are already established. TCP strict can be enabled on a per section basis to turn off mid-session pick-up and enforce the requirement for a three-way handshake. When enabling TCP strict mode for a particular DFW policy, and using a default ANY-ANY Block rule, packets that do not complete the three-way handshake connection requirements and that match a TCP-based rule in this section are dropped. Strict is only applied to stateful TCP rules, and is enabled at the distributed firewall policy level. TCP strict is not enforced for packets that match a default ANY-ANY Allow which has no TCP service specified.
Stateful	A stateful firewall monitors the state of active connections and uses this information to determine which packets to allow through the firewall.
Locked	The policy can be locked to prevent multiple users from editing the same sections. When locking a section, you must include a comment. Some roles such as enterprise administrator have full access credentials, and cannot be locked out. See Role-Based Access Control.

7 Click **Publish**. Multiple policies can be added, and then published together at one time.

The new policy is shown on the screen.

8 Select a policy section and click **Add Rule** and enter a rule name.

9 In the **Sources** column, click the edit icon and select the source of the rule. Groups with Active Directory members can be used for the source text box of an IDFW rule.

See Add a Group.

IPv4, IPv6, and multicast addresses are supported.

Note: IPv6 firewall must have IP Discovery for IPv6 enabled on a connected segment. For more information, see Understanding IP Discovery Segment Profile.

10 In the **Destinations** column, click the edit icon and select the destination of the rule. If not defined, the destination matches any.

See Add a Group.

IPv4, IPv6, and multicast addresses are supported.

11 In the **Services** column, click the edit icon and select services. The service matches **Any** if not defined. See Add a Service.

12 The **Context Profiles** column is not available when adding a rule to the Ethernet category. For all other rule categories, in the **Context Profiles** column, click the edit icon and select a context profile, or click **Add Context Profile**.

See Context Profiles.

Context profiles supports profiles with the attribute type APP ID and Domain (FQDN) Name for use in distributed firewall rules. Multiple App ID context profiles with the attribute type App ID or Domain (FQDN) Name can be used in a distributed firewall rule with services set to **Any**.

Context profiles are not supported when creating IDS rules.

13 Click **Apply** to apply the context profile to the rule.

14 Use **Applied to** to apply the rule to a selected group. When creating a DFW rule using guest introspection, make sure that the **Applied to** field applies to the destination group. By default, the **Applied To** column is set to DFW, and the rule is applied to all workloads. When you change the default, and both the policy level and the rules within have **Applied To** set to **Groups**, then the policy level **Applied To** takes precedence over **Applied To** at the rule level.

Note Groups consisting of only IP addresses, MAC Addresses, or Active Directory groups cannot be used in the **Applied To** text box.

15 In the **Action** column, select an action.

Option	Description
Allow	Allows all L3 or L2 traffic with the specified source, destination, and protocol to pass through the current firewall context. Packets that match the rule, and are accepted, traverse the system as if the firewall is not present.
Drop	Drops packets with the specified source, destination, and protocol. Dropping a packet is a silent action with no notification to the source or destination systems. Dropping the packet causes the connection to be retried until the retry threshold is reached.

Option	Description
Reject	Rejects packets with the specified source, destination, and protocol. Rejecting a packet is a more graceful way to deny a packet, as it sends a destination unreachable message to the sender. If the protocol is TCP, a TCP RST message is sent. ICMP messages with administratively prohibited code are sent for UDP, ICMP, and other IP connections. One benefit of using Reject is that the sending application is notified after only one attempt that the connection cannot be established.
Jump to Application	**Note** This action is only available for the Environment category. Allows the traffic that matches with Environment category rules to continue on for the Application category rules to apply. Use this action when traffic matches with Environment category rules and exits, but you want the Application category rules to apply. For example, if there is an Environment category rule with the action Allow for a specific source and there is an Application category rule with the action Drop for the same source, packets that match the Environment category are allowed through the firewall and further rules are no longer applied. With the Jump to Application action, the packets matches the Environment category rule, but continues on to the Application category rules and the result is that those packets are dropped.

16 Click the status toggle button to enable or disable the rule.

17 Click the gear icon to configure the following rule options:

Option	Description
Logging	Logging is turned off by default. Logs are stored at /var/log/dfwpktlogs.log file on ESXi host.
Direction	Refers to the direction of traffic from the point of view of the destination object. IN means that only traffic to the object is checked, OUT means that only traffic from the object is checked, and In-Out, means that traffic in both directions is checked.
IP Protocol	Enforce the rule based on IPv4, IPv6, or both IPv4-IPv6.
Log Label	Log Label is carried in the Firewall Log when logging is enabled. Only the first 31 characters in the generated log are supported, although you are able to enter a longer label.

18 Click **Publish**. Multiple rules can be added and then published together at one time.

19 The data path realization status of policy with Transport Nodes details shown on the right side of the policy table.

Distributed Firewall Packet Logs

If logging is enabled for firewall rules, you can look at the firewall packet logs to troubleshoot issues.

The log file is `/var/log/dfwpktlogs.log` on ESXi hosts.

Table 16-3. Firewall Log File Variables

Variable	Possible Values
Filter hash	A number that can be used to get the filter name and other information.
AF Value	INET, INET6
Reason	■ match: Packet matches a rule. ■ bad-offset: Datapath internal error while getting packet. ■ fragment: The non-first fragments after they are assembled to the first fragment. ■ short: Packet too short (for example, not even complete to include an IP header, or TCP/UDP header). ■ normalize: Malformed packets that do not have a correct header or a payload. ■ memory: Datapath out of memory. ■ bad-timestamp: Incorrect TCP timestamp. ■ proto-cksum: Bad protocol checksum. ■ state-mismatch: TCP packets that do not pass the TCP state machine check. ■ state-insert: Duplicate connection is found. ■ state-limit: Reached the maximum number of states that a datapath can track. ■ SpoofGuard: Packet dropped by SpoofGuard. ■ TERM: A connection is terminated.
Action	■ PASS: Accept the packet. ■ DROP: Drop the packet. ■ NAT: SNAT rule. ■ NONAT: Matched the SNAT rule, but cannot translate the address. ■ RDR: DNAT rule. ■ NORDR: Matched the DNAT rule, but cannot translate the address. ■ PUNT: Send the packet to a service VM running on the same hypervisor of the current VM. ■ REDIRECT: Send the packet to network service running out of the hypervisor of the current VM. ■ COPY: Accept the packet and make a copy to a service VM running on the same hypervisor of the current VM. ■ REJECT: Reject the packet.
Rule set and rule ID	*rule set/rule ID*
Direction	IN, OUT
Packet length	*length*
Protocol	TCP, UDP, ICMP, or PROTO (protocol number) For TCP connections, the actual reason that a connection is terminated is indicated after the keyword TCP. If TERM is the reason for a TCP session, then an extra explanation appears in the PROTO row. The possible reasons for terminating a TCP connection include: RST (TCP RST packet), FIN (TCP FIN packet), and TIMEOUT (idle for too long) In the example above, it is *RST*. So it means that there is a *RST* packet in the connection that must be reset. For non-TCP connections (UDP, ICMP or other protocols), the reason for terminating a connection is only TIMEOUT.

Table 16-3. Firewall Log File Variables (continued)

Variable	Possible Values
Source IP address and port	*IP address*/*port*
Destination IP address and port	*IP address*/*port*
TCP flags	S (SYN), SA (SYN-ACK), A (ACK), P (PUSH), U (URGENT), F (FIN), R (RESET)
Number of packets	Number of packets. 22/14 - in packets / out packets
Number of bytes	Number of bytes. 7684/1070 - in bytes/ out bytes

The following is a regular log sample for distributed firewall rules:

```
2018-07-03T19:44:09.749Z b6507827 INET match PASS mainrs/1024 IN 52 TCP 192.168.4.3/49627->192.168.4.4/49153 SEW

2018-07-03T19:46:02.338Z 7396c504 INET match DROP mainrs/1024 OUT 52 TCP 192.168.4.3/49676->192.168.4.4/135 SEW

2018-07-06T18:15:49.647Z 028cd586 INET match DROP mainrs/1027 IN 36 PROTO 2 0.0.0.0->224.0.0.1

2018-07-06T18:19:54.764Z 028cd586 INET6 match DROP mainrs/1027 OUT 143 UDP fe80:0:0:0:68c2:8472:2364:9be/546->ff02:0:0:0:0:0:1:2/547
```

The elements of a DFW log file format include the following, separated by a space:

- timestamp:
- last eight digits of the VIF ID of the interface
- INET type (v4 or v6)
- reason (match)
- action (PASS, DROP, REJECT)
- rule set name/ rule ID
- packet direction (IN/OUT)
- packet size
- protocol (TCP, UDP, or PROTO #)
- SVM direction for netx rule hit
- source IP address/source port>destination IP address/destination port
- TCP flags (SEW)

For passed TCP packets there is a termination log when the session has ended:

```
2018-07-03T19:44:30.585Z 7396c504 INET TERM mainrs/1024 OUT TCP RST 192.168.4.3/49627->192.168.4.4/49153 20/16 1718/76308
```

The elements of a TCP termination log include the following, separated by a space:

- timestamp:
- last 8 digits of the VIF ID of the interface
- INET type (v4 or v6)
- action (TERM)
- ruleset name/ rule ID
- packet direction (IN/OUT)
- protocol (TCP, UDP, or PROTO #)
- TCP RST flag
- SVM direction for netx rule hit
- source IP address/source port>destination IP address/destination port
- IN packet count/OUT packet count (all accumulated)
- IN packet size/OUT packet size

The following is a sample of FQDN log file for distributed firewall rules:

```
2019-01-15T00:34:45.903Z 7c607b29 INET match PASS 1031 OUT 48 TCP 10.172.178.226/32808->23.72.199.234/80 S www.sway.com(034fe78d-5857-0680-81e4-d8da6b28d1b4)
```

The elements of an FQDN log include the following, separated by a space:

- timestamp:
- last eight digits of the VIF ID of the interface
- INET type (v4 or v6)
- reason (match)
- action (PASS, DROP, REJECT)
- ruleset name/ rule ID
- packet direction (IN/OUT)
- packet size
- protocol (TCP, UDP, or PROTO #) - for TCP connections, the actual reason that a connection is terminated is indicated after the following IP address
- source IP address/source port>destination IP address/destination port
- TCP flags - S (SYN), SA (SYN-ACK), A (ACK), P (PUSH), U (URGENT), F (FIN), R (RESET

- domain name/UUID where UUID is the binary internal representation of the domain name

The following is a sample of Layer 7 log file for distributed firewall rules:

```
2019-01-15T00:35:07.221Z 82f365ae INET match REJECT 1034 OUT 48 TCP 10.172.179.6/49818->23.214.173.202/80 S APP_HTTP

2019-01-15T00:34:46.486Z 7c607b29 INET match PASS 1030 OUT 48 UDP 10.172.178.226/42035->10.172.40.1/53 APP_DNS
```

The elements of a Layer 7 log include the following, separated by a space:

- timestamp:
- last eight digits of the VIF ID of the interface
- INET type (v4 or v6)
- reason (match)
- action (PASS, DROP, REJECT)
- ruleset name/ rule ID
- packet direction (IN/OUT)
- packet size
- protocol (TCP, UDP, or PROTO #) - for TCP connections, the actual reason that a connection is terminated is indicated after the following IP address
- source IP address/source port>destination IP address/destination port
- TCP flags - S (SYN), SA (SYN-ACK), A (ACK), P (PUSH), U (URGENT), F (FIN), R (RESET
- APP_XXX is the discovered application

Manage a Firewall Exclusion List

Firewall exclusion lists are made of groups that can be excluded from a firewall rule based on group membership.

NSX has system excluded virtual machines, and user excluded groups. In NSX Manager cluster, the first node virtual machine must be manually added to the Distributed Firewall Exclude List. The other NSX Managers, NSX Edge node virtual machines, and Malware Prevention Service virtual machines are automatically added to the read-only the **System Excluded VMs** list. User-defined groups can be excluded from firewall rules, and there are a maximum of 100 groups that can be on the list. IP sets, MAC sets, and Active Directory groups cannot be included as members in a group that is used in a firewall exclusion list.

In NSX 4.0.1.1, exclude lists are supported on a Global Manager (GM) in NSX Federation. On a Local Manager (LM), there will be two exclude lists: one from the GM, and the LM's own exclude list. All members of both lists are excluded.

Antrea groups are not supported in a firewall exclusion list.

VMware recommends excluding the following type of virtual machine/packet processing engines from the distributed firewall using User Excluded Groups:

- Partner Service virtual machines
  - Load balancers
  - Guest Introspection and Service Insertion partner service virtual machines
- Virtual machines that require promiscuous mode

Users should not edit the system generated firewall exclusion list. If edited, traffic may be disrupted.

Procedure

1 Navigate to **Security > Distributed Firewall > Actions > Exclusion List**.

A window appears listing available groups.

2 To view the read-only automated exclusion list, select the **System Excluded VMs** tab. You can filter this list by:

- name
- operating system
- power state
- source
- tag
- tag scope

3 To add a user-defined group to the firewall exclusion list, ensure that you are on the **User Excluded Groups** tab, click the check box next to any group. Then click **Apply**. Note that adding/editing/deleting a group does not change exclusion list membership.

a To create a group, click **Add Group**. See Add a Group.

b To edit a group, click the checkbox next to the group you want to edit, and then click the three dot menu and select **Edit**.

c To delete a group, click the checkbox next to the group you want to delete, and then click the three dot menu and select **Delete**.

d To display group details, click **Expand All**.

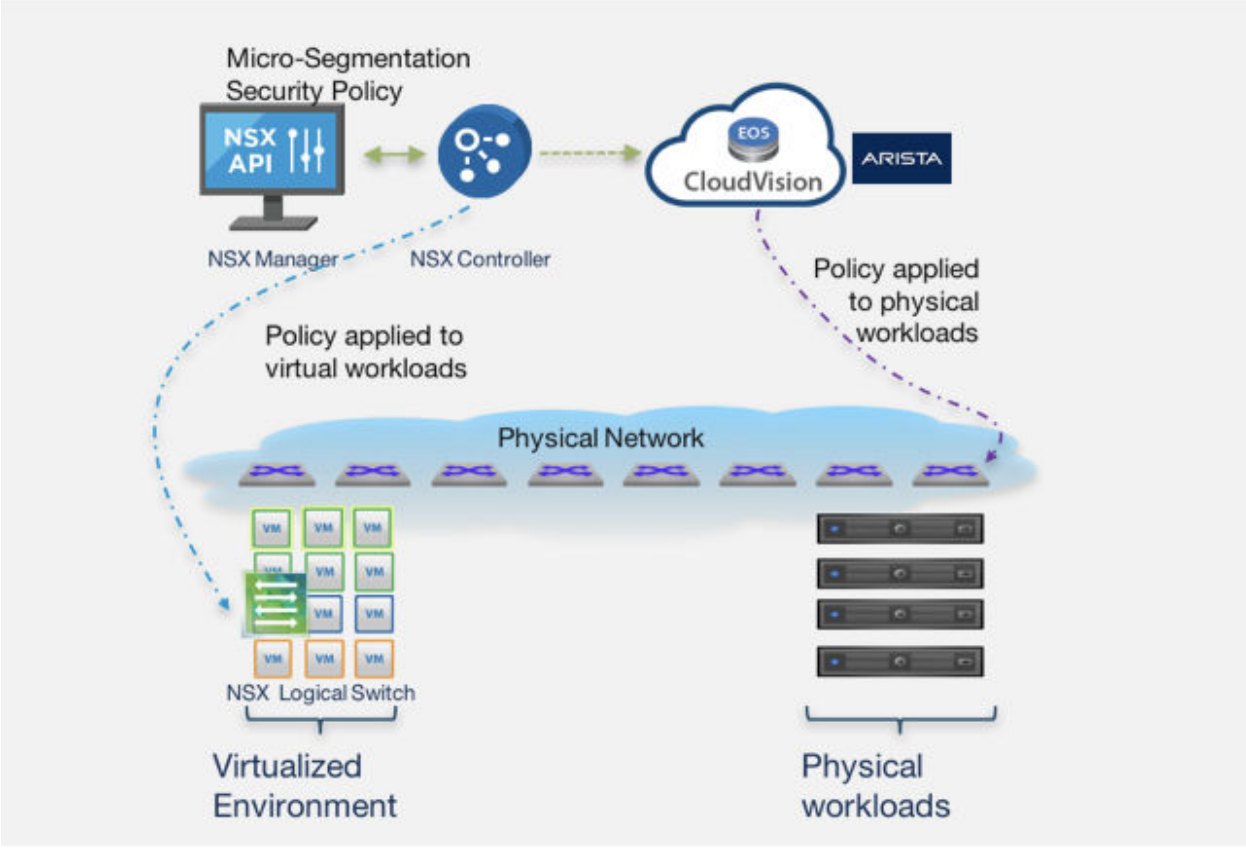
4 Click **Save**.

Extending Security Policies to Physical Workloads

NSX can act as a single point of administration for both virtual and physical workloads.

NSX supports integration with Arista CloudVision eXchange (CVX). This integration facilitates consistent networking and security services across virtual and physical workloads, independent of your application frameworks or physical network infrastructure. NSX does not directly program the physical network switch or router but integrates at the physical SDN controller level, therefore preserving the autonomy of security administrators and physical network administrators.

NSX supports integration with Arista EOS 4.22.1FX-PCS and later.



Limitations

- Arista switches require ARP traffic to exist before firewall rules are applied to an end host that is connected to an Arista switch. Packets can therefore pass through the switch before firewall rules are configured to block traffic.
- Allowed traffic does not resume when a switch crashes or is reloaded. The ARP tables need to be populated again, after the switch comes up, for the firewall rules to be enforced on the switch.
- Firewall rules cannot be applied on the Arista Physical Switch, for FTP passive clients that connect to FTP Server connected to the Arista Physical Switch.
- In CVX HA setup that uses Virtual IP for the CVX cluster, the CVX VM's dvpg's Promiscuous mode, and Forged transmits must be set to Accept. In case they are set to default (Reject), the CVX HA Virtual IP will not be reachable from NSX Manager.

Configure NSX to interact with Arista CVX

Complete the configuration procedure on NSX so that CVX can be added as an enforcement point in NSX and NSX can interact with CVX.

Prerequisites

Obtain the virtual IP address for the Arista CVX cluster.

Procedure

1 Log in to NSX Manager as a root user and run the following command to retrieve the thumbprint for CVX:

```
openssl s_client -connect <virtual IP address of CVX cluster> | openssl x509 -noout
-fingerprint -sha256
```

Sample output:

```
depth=0 CN = self.signed
verify error:num=18:self signed certificate
verify return:1
depth=0 CN = self.signed
verify return:1
SHA256
Fingerprint=35:C1:42:BC:7A:2A:57:46:E8:72:F4:C8:B8:31:E3:13:5F:41:95:EF:D8:1E:E9:3D:F0:CC:3
B:09:A2:FE:22:DE
```

2 Edit the retrieved thumbprint to use only lower case characters and exclude any colons in the thumbprint.

Sample of edited thumbprint for CVX:

```
35c142bc7a2a5746e872f4c8b831e3135f4195efd81ee93df0cc3b09a2fe22de
```

3 Call the `PATCH /policy/api/v1/infra/sites/default/enforcement-points` API and use the CVX thumbprint to create an enforcement endpoint for CVX. For example:

```
PATCH https://<nsx-manager>/policy/api/v1/infra/sites/default/enforcement-points/cvx-
default-ep
{
"auto_enforce": "false",
"connection_info": {
"enforcement_point_address": "<IP address of CVX>",
"resource_type": "CvxConnectionInfo",
"username": "cvpadmin",
"password": "1q2w3e4rT",
"thumbprint": "65a9785e88b784f54269e908175ada662be55f156a2dc5f3a1b0c339cea5e343"
}
}
```

4 Call the `GET /policy/api/v1/infra/sites/default/enforcement-points` API to retrieve the endpoint information. For example:

```
https://<nsx-manager>/policy/api/v1/infra/sites/default/enforcement-points/cvx-default-ep
{
"auto_enforce": "false",
"connection_info": {
"enforcement_point_address": "<IP address of CVX>",
"resource_type": "CvxConnectionInfo",
"username": "admin",
"password": "1q2w3e4rT",
"thumbprint": "35c142bc7a2a5746e872f4c8b831e3135f4195efd81ee93df0cc3b09a2fe22de"
}
}
```

Sample output:

```
{
  "connection_info": {
    "thumbprint": "35c142bc7a2a5746e872f4c8b831e3135f4195efd81ee93df0cc3b09a2fe22de",
    "enforcement_point_address": "192.168.2.198",
    "resource_type": "CvxConnectionInfo"
  },
  "auto_enforce": false,
  "resource_type": "EnforcementPoint",
  "id": "cvx-default-ep",
  "display_name": "cvx-default-ep",
  "path": "/infra/sites/default/enforcement-points/cvx-default-ep",
  "relative_path": "cvx-default-ep",
  "parent_path": "/infra/sites/default",
  "marked_for_delete": false,
  "_system_owned": false,
  "_create_user": "admin",
  "_create_time": 1564036461953,
  "_last_modified_user": "admin",
  "_last_modified_time": 1564036461953,
  "_protection": "NOT_PROTECTED",
  "_revision": 0
}
```

5 Call the `POST /api/v1/notification-watchers/` API and use the CVX thumbprint to create a notification ID. For example:

```
POST https://<nsx-manager>/api/v1/notification-watchers/
{
"server": "<virtual IP address of CVX cluster>",
"method": "POST",
"uri": "/pcs/v1/nsgroup/notification",
"use_https": true,
"certificate_sha256_thumbprint":
"35c142bc7a2a5746e872f4c8b831e3135f4195efd81ee93df0cc3b09a2fe22de",
"authentication_scheme": {
"scheme_name":"BASIC_AUTH",
```

```
"username":"cvpadmin",
"password":"1q2w3e4rT"
}
}
```

6 Call the `GET /api/v1/notification-watchers/` to retrieve the notification ID.

Sample output:

```
{
  "id": "a0286cb6-de4d-41de-99a0-294465345b80",
  "server": "192.168.2.198",
  "port": 443,
  "use_https": true,
  "certificate_sha256_thumbprint":
"35c142bc7a2a5746e872f4c8b831e3135f4195efd81ee93df0cc3b09a2fe22de",
  "method": "POST",
  "uri": "/pcs/v1/nsgroup/notification",
  "authentication_scheme": {
    "scheme_name": "BASIC_AUTH",
    "username": "cvpadmin"
  },
  "send_timeout": 30,
  "max_send_uri_count": 5000,
  "resource_type": "NotificationWatcher",
  "display_name": "a0286cb6-de4d-41de-99a0-294465345b80",
  "_create_user": "admin",
  "_create_time": 1564038044780,
  "_last_modified_user": "admin",
  "_last_modified_time": 1564038044780,
  "_system_owned": false,
  "_protection": "NOT_PROTECTED",
  "_revision": 0
}
```

7 Call the `PATCH /policy/api/v1/infra/domains/default/domain-deployment-maps/cvx-default-dmap` API to create a CVX domain deployment map. For example:

```
PATCH https://<nsx-manager>/policy/api/v1/infra/domains/default/domain-deployment-maps/cvx-
default-dmap
{

"display_name": "cvx-deployment-map",

"id": "cvx-default-dmap",

"enforcement_point_path": "/infra/sites/default/enforcement-points/cvx-default-ep"

}
```

8 Call the `GET /policy/api/v1/infra/domains/default/domain-deployment-maps` API to retrieve the deployment map information.

Configure Arista CVX to interact with NSX

After configuring NSX, complete the configuration procedure on Arista CloudVision eXchange (CVX) to enable CVX to interact with NSX.

Prerequisites

NSX has registered the CVX as an enforcement point.

Procedure

1 Log in to NSX Manager as a root user and run the following command to create a thumbprint for CVX to communicate with NSX Manager:

```
openssl s_client -connect <IP address of nsx-manager>:443 | openssl x509 -pubkey -noout |
openssl rsa -pubin -outform der | openssl dgst -sha256 -binary | openssl base64
```

Sample output:

```
depth=0 C = US, ST = CA, L = Palo Alto, O = VMware Inc., OU = NSX, CN = nsx-mgr
verify error:num=18:self signed certificate
verify return:1
depth=0 C = US, ST = CA, L = Palo Alto, O = VMware Inc., OU = NSX, CN = nsx-mgr
verify return:1
writing RSA key
S+zwADluzeNf+dnffDpYvgs4YrS6QBgyeDry40bPgms=
```

2 Run the following commands from the CVX CLI:

```
cvx
no shutdown
service pcs
no shutdown
controller <IP address of nsx-manager>
username <NSX administrator user name>
password <NSX administrator password>
enforcement-point cvx-default-ep
pinned-public-key <thumbprint for CVX to communicate with NSX
                   Manager>
notification-id <notification ID created while registering CVX with NSX>
end
```

3 Run the following command from the CVX CLI to check the configuration:

```
show running-config
```

Sample ouput:

```
cvx
      no shutdown
   source-interface Management1
   !
   service hsc
```

```
      no shutdown

   !
   service pcs
      no shutdown
      controller 192.168.2.80
      username admin
      password 7 046D26110E33491F482F2800131909556B
      enforcement-point cvx-default-ep
      pinned-public-key sha256//S+zwADluzeNf+dnffDpYvgs4YrS6QBgyeDry40bPgms=
      notification-id a0286cb6-de4d-41de-99a0-294465345b80
```

4 Configure `tag` on the ethernet interface of the physical switch that connects to the physical server. Run the following commands on the physical switch managed by CVX.

```
configure terminal
interface ethernet 4
tag phy_app_server
end
copy running-config startup-config
Copy completed successfully.
```

5 Run the following command to verify tag configuration for the switch:

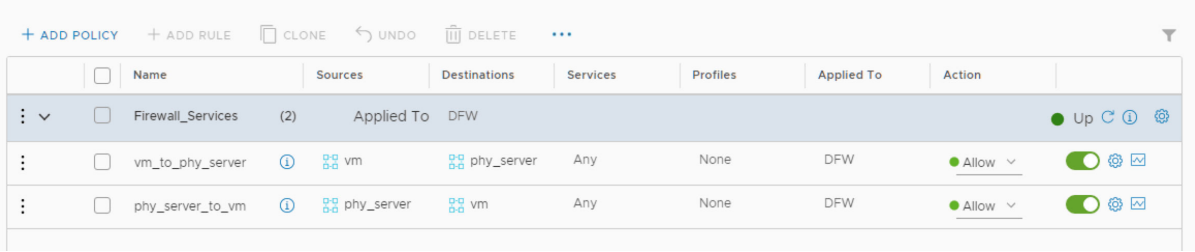
```
show running-config section tag
```

Sample output:

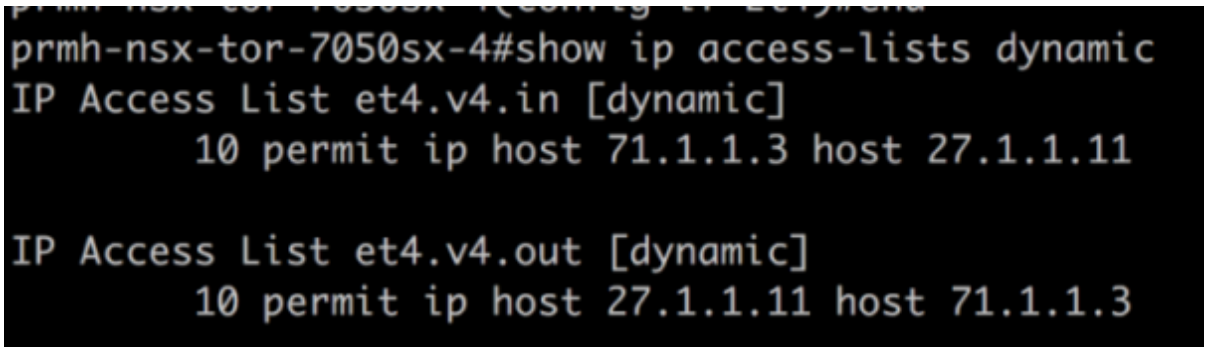
```
interface Ethernet4
   description connected-to-7150s-3
   switchport trunk allowed vlan 1-4093
   switchport mode trunk
   tag sx4_app_server
```

IP addresses that are learnt on the tagged interfaces, using ARP, are shared with NSX.

6 Log in to NSX Manager to create and publish firewall rules for the physical workloads managed by CVX. See Chapter 16 Security for more information on creating rules. For example:



NSX policies and rules published in NSX appear as dynamic ACLs on the physical switch managed by CVX.



For more information, see CVX HA set up, CVX HA Virtual IP setup, and Physical Switch Mlag Setup

Shared Address Sets

Security groups based on dynamic or logical objects can be created and used in the **Applied to** text box of distributed firewall rules.

Because address sets are dynamically populated based on virtual machine name or tags, and must be updated on each filter, they can exhaust the available amount of heap memory on hosts to store DFW rules and IP address sets.

In NSX version 2.5 and later, a feature called Global or Shared Address Sets, makes address sets shared across all the filters. While each filter can have different rules, based on **Applied To**, the address sets members are constant across all the filters. This feature is enabled by default, reducing heap memory use. It cannot be disabled.

In NSX version 2.4 and earlier, Global or Shared Address Sets is disabled, and environments with heavy distributed firewall rules might experience VSIP heap exhaustion.

Export or Import a Firewall Configuration

For a consolidated view of your policy sections and rules, you can export your firewall configuration to a file. NSX creates a report of your firewall configuration as a CSV file. You can also import a firewall configuration and view it as a draft in NSX.

- You can choose to export your published configuration, an auto-saved draft, or a manually saved draft.
- When exporting a firewall configuration on a Local Manager appliance, only the Local Manager configuration is exported. Configuration that has been synced from a Global Manager is not exported as part of the Local Manager configuration.
- The import operation is not available on a Global Manager appliance.
- The exported CSV file and the metadata file together are available for download as a ZIP file.
- You can only run one export operation and one import operation simultaneously. Running multiple export or multiple import operations simultaneously is not supported.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Security > Distributed Firewall**.

3 Export the firewall configuration.

- To export the current configuration, select **Actions > Export FW Configuration**.
- To export an auto-saved draft or a manual draft, select **Actions > View** and click the name of the draft configuration to open **Draft** details. Click **Export Draft**.

  The **View Draft Details** window displays any differences that exist between the saved configuration and the last published configuration.

4 Enter a passphrase and click **Export**.

A notification is displayed when the configuration has been exported.

You cannot publish a configuration when the export operation is in progress. If necessary, you can cancel the export operation.

5 Click **Download** to save the ZIP file containing the CSV and metadata files.

6 To import a firewall configuration, navigate to **Security > Distributed Firewall**, and select **Actions > Import**.

When importing rules with groups, the groups must be created on the destination environment without typos. If not, you will get a `Deleted_Object` error message instead of the group name when importing the rules.

Editing the name of the Group to fix the typo does not fix the issue, because the UUID stays with the the original name.

7 Browse to select the ZIP file containing the configuration that you want to import. Enter a name and the passphrase that was used when saving the configuration, and click **Import**.

Ensure that you select a ZIP file that has not been modified after it was downloaded.

A notification is displayed if the file to be imported is corrupt or the incorrect passphrase has been used.

Note You can only import configuration that has been defined in NSX. Importing third-party firewall configuration is not supported.

The imported configuration is saved as a manual draft in NSX. You can edit the draft and then publish it as required. For more information on working with drafts, see Firewall Drafts.

Gateway Firewall

Gateway firewall represents rules applied at the perimeter firewall.

There are predefined categories under the **All Shared Rules** view, where rules across all gateways are visible. Categories are evaluated from left to right, and the rules within the category are evaluated top down. The category names can be changed using the API.

Emergency Used for Quarantine. Can also be used for Allow rules.	**System** These rules are automatically generated by NSX and are specific to internal control plane traffic, such as, BFD rules, VPN rules, and so on. **Note** Do not edit System rules.	**Shared Pre Rules** These rules are globally applied across gateways.	**Local Gateway** These rules are specific to a particular gateway.	**Auto Service Rules** These are auto-plumbed rules applied to the data plane. You can edit these rules as required.	**Default** These rules define the default gateway firewall behavior.

Supported Gateway Firewall Features on NSX Edge

Before you configure Gateway Firewall features, make sure that the NSX Edge form factor deployed in your environment supports the features.

Table 16-4. Gateway Firewall features supported on NSX Edge form factor

Features/NSX Edge Form Factor	Small 2 vCPU, 4GB RAM (POC only)	Medium 4 vCPU, 8 GB RAM	Large 8 vCPU, 32 GB RAM	Extra Large 16 vCPU, 64 GB RAM	Bare Metal
L3-L4 Firewall	Yes	Yes	Yes	Yes	Yes
User ID-based Access Control	Yes	Yes	Yes	Yes	Yes
Application Access Control	No	Yes	Yes	Yes	Yes
URL Filtering	No	Yes	Yes	Yes	Yes
FQDN Analysis	No	Yes	Yes	Yes	Yes
IDPS	No	No	Yes	Yes	Yes
Malware Detection	No	No	No	Yes	No
Sandboxing for unknown Threats	No	No	No	Yes	No
TLS Inspection	No	No	Yes	Yes	Yes

Table 16-4. Gateway Firewall features supported on NSX Edge form factor (continued)

Features/NSX Edge Form Factor	Small 2 vCPU, 4GB RAM (POC only)	Medium 4 vCPU, 8 GB RAM	Large 8 vCPU, 32 GB RAM	Extra Large 16 vCPU, 64 GB RAM	Bare Metal
L2 and L3 VPN	Yes	Yes	Yes	Yes	Yes
Static, Dynamic Routing	Yes	Yes	Yes	Yes	Yes

Gateway Firewall Settings

Gateway Firewall Settings include options for gateway specific settings, FQDN analysis, and URL filtering.

Gateway Specific Settings

To view Gateway Firewall Settings navigate to **Security > Gateway Firewall > Settings > Gateway Specific Settings**. Click **TURN ON** for the Tier-1 or Tier-0 gateway firewall you want to activate.

FQDN Analysis

Navigate to **Security > Gateway Firewall > Settings > FQDN Analysis**. Select the edge cluster/ node name upon which you want to activate FQDN Analysis. Click **TURN ON**. Once activated, the URL database will be downloaded to each cluster member. FQDN Analysis is applied to all internet-bound traffic passing through the edge cluster.

URL Filtering

Navigate to **Security > Gateway Firewall > Settings > URL Filtering** to view the **Reject with Response** message.

The **Reject with Response** page is sent only for http traffic. The response page will contain the URL (first 10 bytes show), Category, Source IP and message-text. Enter the message for the **Reject with Response** page. Click **Preview Page** to view the page that will be sent when access to a URL is blocked by a policy.

Add a Gateway Firewall Policy and Rule

To enable identity firewall, navigate to **Security > Gateway Firewall**, and click **Actions**, **General Settings**. Toggle the bar to enable identity firewall, noting that you must have active directory.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Select **Security > Gateway Firewall**.
3. To enable Gateway Firewall select **Actions > General Settings**, and toggle the status button. Click **Save**.

4 Click **Add Policy**, for more about categories see Gateway Firewall.

5 Enter a **Name** for the new policy section.

6 Select the policy **Destination**.

7 Click the gear icon to configure the following policy settings:

Settings	Description
TCP Strict	A TCP connection begins with a three-way handshake (SYN, SYN-ACK, ACK), and typically ends with a two-way exchange (FIN, ACK). In certain circumstances, the firewall may not see the three-way handshake for a particular flow (i.e. due to asymmetric traffic). By default, the firewall does not enforce the need to see a three-way handshake, and will pick-up sessions that are already established. TCP strict can be enabled on a per section basis to turn off mid-session pick-up, and enforce the requirement for a three-way handshake. When enabling TCP strict mode for a particular firewall policy and using a default ANY-ANY Block rule, packets that do not complete the three-way handshake connection requirements and that match a TCP-based rule in this policy section are dropped. Strict is only applied to stateful TCP rules, and is enabled at the gateway firewall policy level. TCP strict is not enforced for packets that match a default ANY-ANY Allow which has no TCP service specified.
Stateful	A stateful firewall monitors the state of active connections, and uses this information to determine which packets to allow through the firewall.
Locked	The policy can be locked to prevent multiple users from making changes to the same sections. When locking a section, you must include a comment.

8 Click **Publish**.

Multiple Policies can be added, and then published together at one time.

The new policy is shown on the screen.

9 Select a policy section and click **Add Rule**.

10 Enter a name for the rule. IPv4, and IPv6 addresses are supported.

11 In the **Sources** column, click the edit icon and select the source of the rule. Groups with Active Directory members can be used for the source box of an IDFW rule. See Add a Group.

12 In the **Destinations** column, click the edit icon and select the destination of the rule. If not defined, the destination matches any. See Add a Group.

13 In the **Services** column, click the pencil icon and select services. The service matches any if not defined. See Add a Service.

14 For Tier-1 gateways, in the **Profiles** column, click the edit icon and select a context profile, or L7 Access Profile. Or, create new profiles. See Profiles.

- A security rule can contain either a context profile or an L7 access profile, but not both.
- Context profiles and L7 access profiles are not supported on tier-0 gateway firewall policy.
- Gateway firewall rules do not support context profiles with attribute type Domain(FQDN) Name.
- Gateway firewall rules support L7 access profiles with attribute type App ID, URL Category, Custom URL and URL Reputation. The attribute type App ID supports multiple sub attributes.

Multiple App ID context profiles can be used in a firewall rule with services set to **Any**. Only a single L7 Access profile can be used within a single gateway firewall rule.

15 Click **Apply**.

16 Click the pencil icon for the **Applied To** column to change the scope of enforcement per rule. From the **Applied To | New Rule** dialog box, click the **Categories** drop-down menu to filter by object type such as interfaces, labels, and VTIs to select those specific objects.

By default, gateway firewall rules are applied to all the available uplinks and service interfaces on a selected gateway.

For URL filtering, **Applied To** can only be Tier-1 gateways.

17 In the **Action** column, select an action.

Option	Description
Allow	Allows all traffic with the specified source, destination, and protocol to pass through the current firewall context. Packets that match the rule, and are accepted, traverse the system as if the firewall is not present. The rule action with an L7 access profile must be **Allow.**
Drop	Drops packets with the specified source, destination, and protocol. Dropping a packet is a silent action with no notification to the source or destination systems. Dropping the packet causes the connection to be retried until the retry threshold is reached.
Reject	Rejects packets with the specified source, destination, and protocol. Rejecting a packet sends a destination unreachable message to the sender. If the protocol is TCP, a TCP RST message is sent. ICMP messages with administratively prohibited code are sent for UDP, ICMP, and other IP connections. The sending application is notified after one attempt that the connection cannot be established.

18 Click the status toggle button to activate or deactivate the rule.

19 Click the gear icon to set logging, direction, IP protocol, and comments.

Option	Description
Logging	Logging can be turned off or on. Logs are stored at /var/log/syslog on the Edge.
Direction	The options are **`In`**, **`Out`**, and **`In/Out`**. The default is **`In/Out`**. This field refers to the direction of traffic from the point of view of the destination object. **`In`** means that only traffic to the object is checked, **`Out`** means that only traffic from the object is checked, and **`In/Out`** means that traffic in both directions is checked.
IP Protocol	The options are **`IPv4`**, **`IPv6`**, and **`IPv4_IPv6`**. The default is **`IPv4_IPv6`**.

Note Click the graph icon to view the flow statistics of the firewall rule. You can see information such as the byte, packet count, and sessions.

20 Click **Publish**. Multiple rules can be added and then published together at one time.

21 On each policy section, click the **Info** icon to view the current status of edge firewall rules that are pushed to edge nodes. Any alarms generated when rules were pushed to edge nodes are also displayed.

22 To view consolidated status of policy rules that are applied to edge nodes, make the API call.

```
GET https://<policy-mgr>/policy/api/v1/infra/
realized-state/status?intent_path=/infra/domains/default/gateway-policies/
<GatewayPolicy_ID>&include_enforced_status=true
```

TLS Inspection

TLS Inspection is used to detect and prevent advanced threats over encrypted TLS channels. TLS Inspection transparently decrypts encrypted traffic and makes it available for advanced security features such as IDS/IPS, Malware Prevention, and URL Filtering. This provides visibility into the encrypted traffic without offloading and while retaining end-to-end encryption.

Without TLS Inspection, even if you enable all the advanced security features for the gateway firewall you cannot enforce or have visibility into the encrypted traffic that may have hidden malware inside the packets. TLS decryption allows administrators to have more effective access control and threat detection and prevention in encrypted traffic.

TLS Inspection Support

This topic describes support for TLS Inspection in NSX.

TLS Inspection support includes:

- Support on tier-1 gateways only.
- Support for TLS version 1.0 through 1.2 and. TLS 1.2 with Perfect Forward Secrecy (PFS). If version 1.3 is used, the NSX proxy negotiates to an earlier version and establishes a connection.

- Leverages TLS Server Name Indication (SNI) in TLS client hello to classify the traffic.
- Visibility into encrypted traffic without offloading while retaining end-to-end encryption.
- TLS decryption on gateway firewalls to intercept the traffic and decrypt it to feed to the advanced firewall security features.
- TLS Inspection policies to create a set of rules that describe conditions to match and perform a predefined action.
- The TLS Inspection policy rules support the bypass, external, and internal decryption action profiles.

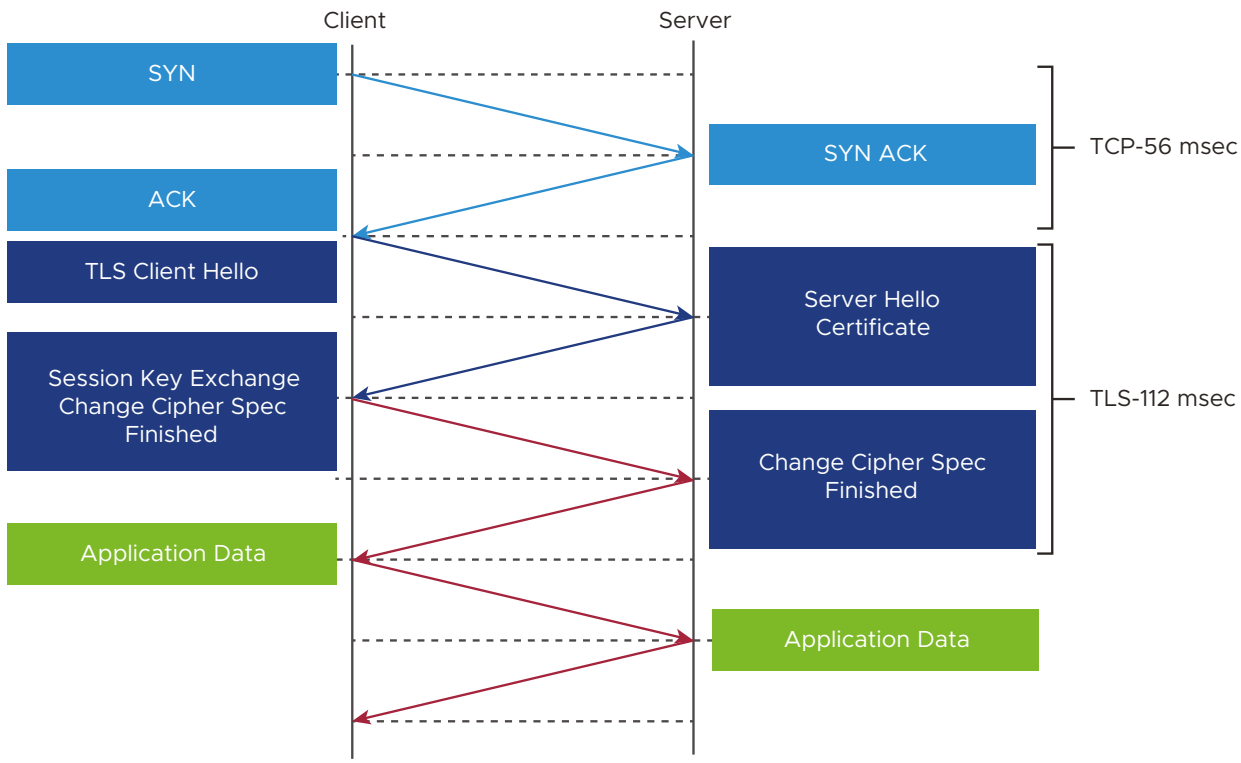
TLS Inspection Concepts

TLS Inspection detects and prevents advanced threats in your network over encrypted TLS channels. This topic includes concepts associated with TLS Inspection features.

TLS Protocol

This topic describes how TLS protocol handshake works to establish an encrypted channel between the client and the server. The following TLS protocol illustration provides the various steps involved to form an encrypted channel.

Figure 16-3. TLS Protocol



To summarize the TLS protocol:

- TLS initiates a TLS session over an established TCP session between the client and the server (aka a three-way handshake).
- The client sends a Client Hello that includes the supported TLS version and cipher and the Server Name Indication (SNI) extension. SNI in the TLS Client Hello is what TLS Inspection uses to classify the traffic using the context profile to use the internal, external, or bypass decryption profiles.
- The server responds with the server certificate for authentication and identification and a Server Hello with the version and cipher proposed by the client.
- Once the client validates the certificate and verifies the final version and cipher, it generates a symmetric session key and sends it to the server.
- To initiate the secure TLS tunnel which exchanges application data over the encrypted TLS channel, the server validates the session key and sends the finished message.

By default the TLS protocol only proves the identity of the server to the client using X.509 certificate and the authentication of the client to the server is left to the application layer.

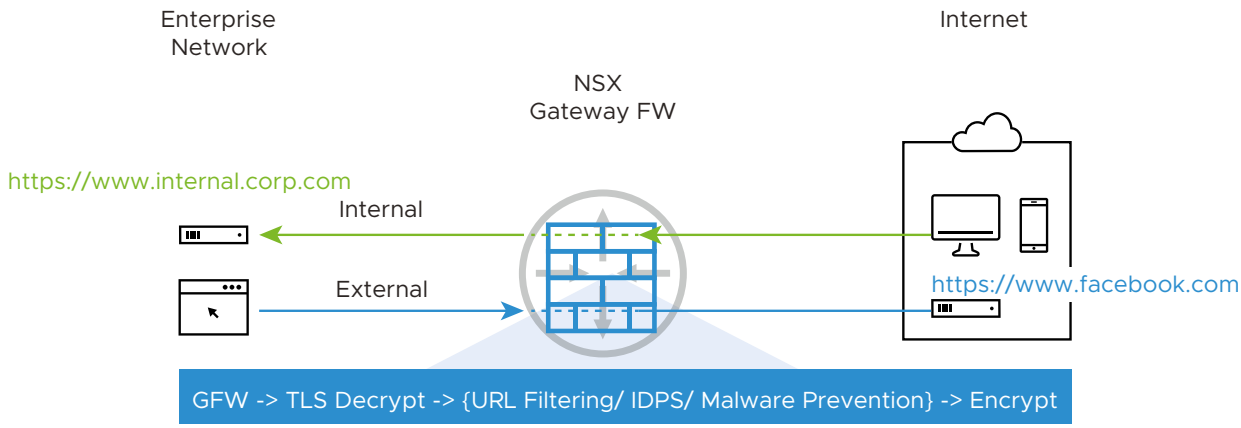
TLS Decryption Types

The TLS Inspection feature allows users to define policies to decrypt or to bypass the decryption, the TLS Inspection feature allows two types of decryption:

- Internal TLS Decryption - for traffic going to an Enterprise internal service where you own the service, certificate, and the private key. This is also called TLS reverse-proxy or inbound decryption.
- External TLS Decryption - for traffic going to an external service (Internet) where the Enterprise does not own the service, its certificate, and the private key. This is also called TLS forward proxy or outbound decryption.

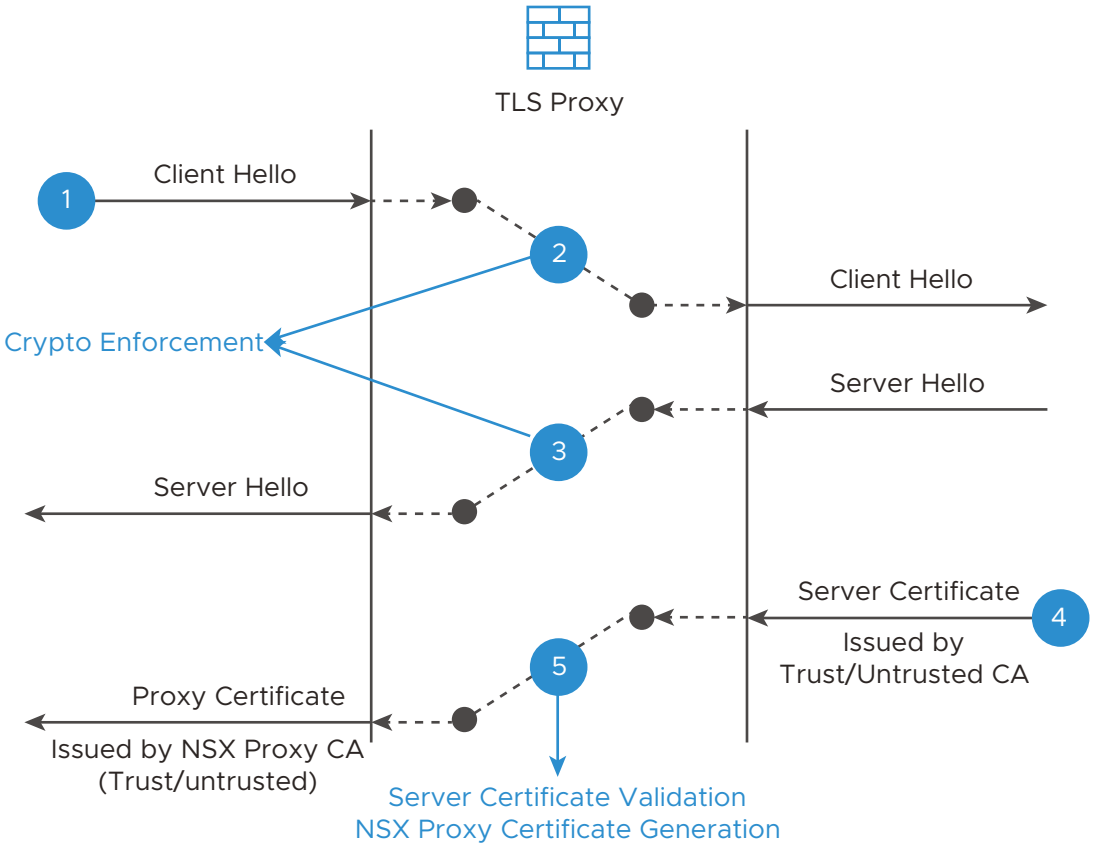
The following diagram depicts how the traffic is handled by the TLS internal and external decryption types.

Figure 16-4. NSX TLS Decryption Types



The following diagram and table explain how NSX TLS External decryption works with NSX.

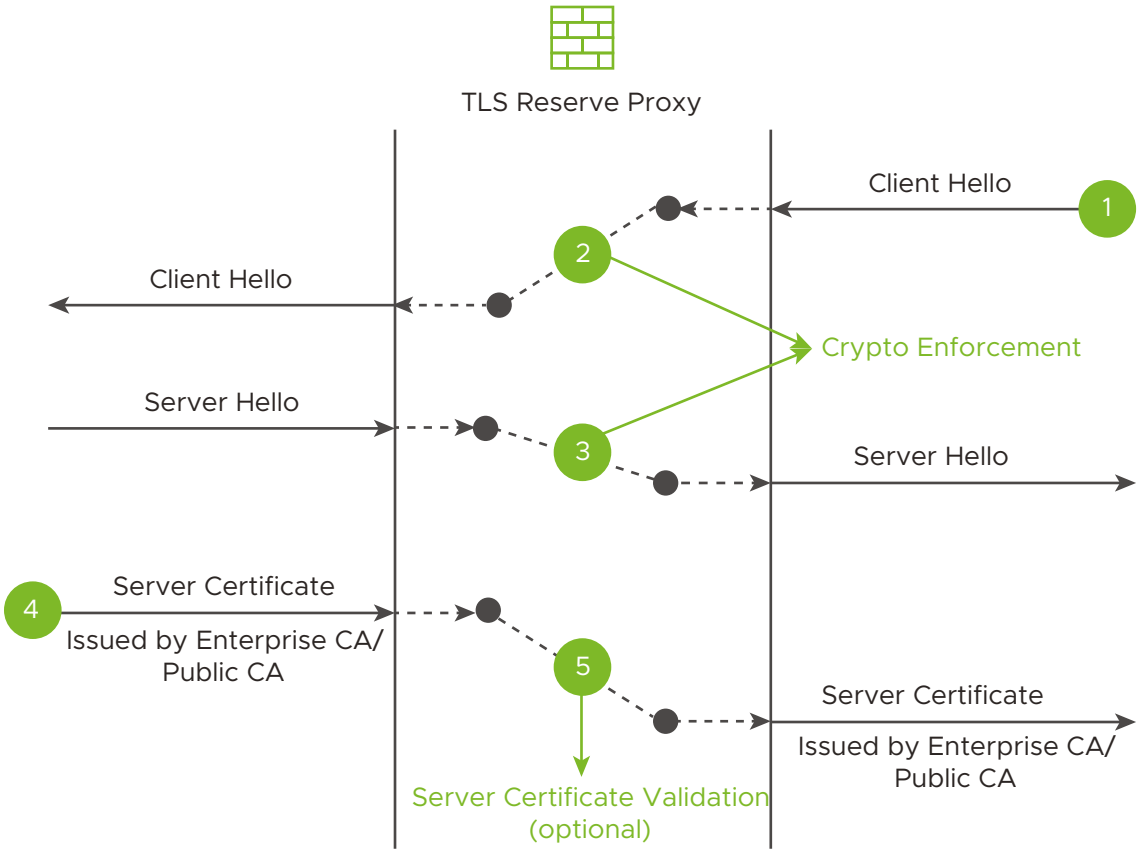
Figure 16-5. How External Decryption Works



Callout	Workflow
1	TLS client hello SNI matches the TLS Inspection policy context profile.
2	NSX intercepts the TLS session from the client and initiates a new session to the intended server.
3	NSX enforces TLS version and cipher (which is configurable). .
4	The server responds to the client with a TLS certificate
5	NSX validates the server certificate using the trusted CA bundle and generates a proxy CA certificate dynamically and presents that to the client.

The following diagram and table explain how NSX TLS Internal decryption works with NSX.

Figure 16-6. How Internal Decryption Works



Callout	Workflow
1	TLS client hello SNI matches the TLS Inspection policy context profile configured for internal domain.
2	NSX intercepts the TLS session from the client and initiates a new session to intended server.
3	NSX enforces TLS version/cipher (configurable).
4	Server responds with certificate as part of TLS handshake (validation optional).
5	NSX presents the certificate of the server, which was uploaded as part of the configuration, to the client.

TLS Inspection Policies

A TLS Inspection policy applies to the selected tier-1 gateway firewall or firewalls. The first time you add a TLS Inspection policy, you can use the wizard, or you can manually configure the policy and associated rules. This topic describes the concepts and creation of TLS Inspection policies.

TLS Inspection Policy Categories

NSX TLS inspection provides the following three categories for easy policy management. Similar to gateway firewall categories, you can use any of the categories based on requirement to define TLS inspection policies.

- Pre-Rules - Defines the policy for multiple gateways.
- Local Gateway - Defines specific policies.
- Default (post-rules) - This TLS Default category is different than the gateway policy rules as it does not contain any out-of-the-box rule or policy default. It also allows you to define post rules in the Default category (which is not available in the gateway firewall table). For example, the use case might be to add some common policies to multiple gateways after the local gateway configuration.

Create a TLS Inspection Policy

To simplify the configuration of the first TLS Inspection policy, you can use the TLS Inspection wizard or manually create your policy using the UI. This topic does not describe the wizard configuration, only the manual configuration steps.

The wizard provides a walk-through of the TLS Inspection configuration workflow for your tier-1 gateway firewalls. The wizard displays on the TLS Inspection home page only for the first policy, but you can access the wizard in the All Shared Rules and Gateway Specific Rules tabs. You can skip the configuration wizard and complete the policy creation and the decryption action profile setup manually by clicking **Skip** on the opening page.

Prerequisites

These prerequisites are valid for TLS Inspection in policies.

Activate the following settings. By default, they are deactivated.

- Activating TLS Inspection settings per gateway.

  Navigate to **Security > TLS Inspection** and select the **Settings** tab. Select a gateway or gateways from the list of TLS-enabled gateways and click **Turn On**.

- Activating URL Database on the Edge cluster.

  Navigate to **Security > General Settings > URL Database**. Edge nodes must have Internet connectivity so the NSX Threat Intelligence Cloud Service (NTICS) can complete URL database downloads.

- To view TLS Inspection statistics using the Security dashboard, deploy NSX Application Platform on your NSX 3.2 or later environment and ensure it is in a good state. A specific license is required for time-series monitoring. For details, see the *Deploying and Managing NSX Application Platform* guide and Monitoring Security Statistics.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Security > TLS Inspection**.

3 Select the category to define the policy, then click **Add Policy**.

4 Enter a name for the new policy.

5 (Optional) If you want to prevent multiple users from making changes to the section, click the **Advanced Configuration** icon, then click **Locked** and **Apply**.

6 Select the policy you created, then click **Add Rule**.

Variable	Description
Source, Destination, and L4 services	Matches the same fields of the traffic coming in as the gateway firewall rule.
Context profile	Define and select context profile for classifying the traffic based on URL Category, Reputation, and Domain name. For details, see Context Profiles.
Decryption action profile	Define and select the decryption profile for the matched traffic. This could be external, internal, and bypass profiles. For details, see Creating TLS Decryption Action Profiles.
Applied to	Select one or more tier-1 gateways.

7 Click **Publish**.

You have completed your policy creation.

Creating TLS Decryption Action Profiles

You can create three types of decryption action profiles in TLS Inspection. This topic describes the profiles and how to use them.

The decryption action profiles are:

- Bypass decryption profiles - This profile is used to bypass the decryption of traffic destined for specific categories of websites, for compliance and privacy reasons. For example, healthcare and financial websites. The bypass profile cannot be changed.
- External decryption profiles - These profiles are for TLS connections destined to a service not owned by the Enterprise. For example, http://merriam-webster.com. In general, Internet websites.
- Internal decryption profiles - These profiles are for TLS connections destined to a service owned by the Enterprise itself. For example, http://www.corp.internal.com.

Create an External Decryption Profile

This topic provides the steps to configure an external decryption action profile manually.

Prerequisites

- Have the correct user role and permissions to set up TLS Inspection.

- Have a Trusted Proxy CA and Untrusted Proxy CA certificate imported or ready to be imported or have the related information to generate a certificate.

Procedure

1 With admin privileges, login into NSX Manager.

2 Navigate to **Security > TLS Inspection > Profiles**.

3 Click **Add Decryption Action Profile > External Decryption**.

4 Enter a name for the new profile.

5 (Optional) Select a profile setting: Balanced (default), High Fidelity, High Security, or use Custom to change the sub-settings.

Profile Setting	Description
Invalid Certificates: Allow or Block & Log	Set rules to allow or block traffic when an invalid certificate is presented by server. If Allow is selected and the server presents with an expired or untrusted certificate, this choice allows the connection to proceed by sending an untrusted proxy certificate to the client.
Decryption Failure: Bypass & Log or Block & Log	Sets what to do when there is decryption failure which could be due to mTLS (mutual TLS) or certificate pinning in use. If Bypass & Log is selected, then NSX caches this domain, and all subsequent connections to the domain are bypassed.
Crypto Enforcement: Transparent or Enforce	Sets the minimum and maximum TLS versions and cipher suites for the client and server. You can bypass this using the Transparent option

6 (Optional) Modify Idle connection timeout. This is the time in seconds the server can remain idle after establishing a TCP connection. Default is 5400 seconds. Keep this timeout lower than the gateway firewall idle timeout settings.

7 (Optional) Select Trusted CA settings to select Trusted CA Bundle, CRLs, and the OCSP stapling option.

Option	Description
Trusted CA Bundle	Validates the certificate that the external service presents to NSX. You can use the default trusted CA bundle or import a new CA bundle, then choose multiple bundles per profile if needed. This bundle is not automatically updated, so you must update it as necessary. For more details, see Import or Update a Trusted CA Bundle under Certificate Management.
CRLs	NSX also includes a CRL (Certificate revocation list) to validate the server presented certificate. You can use the default CRL or import a new CRL, then choose multiple CRLs per profile if needed. This CRL is not automatically updated, so you must update it as necessary. For more details, see Importing and Retrieving CRLs under Certificate Management.
Require OCSP Stapling	To enforce OSCP stapling for the presented server certificate. In OCSP stapling, the server that owns the certificate queries OCSP responder and includes the received OCSP timestamped and signed response as CertificateStatusRequest extension along with its certificate. If the server has a chained certificate, then the server must do OCSP Stapling for all the intermediate CA certs as well.

8 To import or generate a trusted or untrusted proxy CA, select the Proxy CA dropdown, select the **Trusted Proxy CA** or the **Untrusted Proxy CA tab**, then do one of the following:

- Select **Import > CA Certificate**.
- Select **Generate > Self Signed CA Certificate**.

  Enter the required details and click Save. For details about importing proxy CAs, see Importing and Replacing Certificates.

9 To save the profile, which can then be used for TLS inspection policies, select **Save**.

Results

You are now able to use the decryption action profile to set up external decryption rules on your tier-1 gateways.

What to do next

Create TLS Inspection external decryption policies and rules.

Create an Internal Decryption Action Profile

This topic provides the steps to configure an internal decryption action profile manually.

Prerequisites

- Have the correct user role and permissions to set up TLS Inspection.
- Have an internal server certificate imported or ready to be imported or have the related information to generate the certificate.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Security > TLS Inspection > Profiles**.

3 Click **Add Decryption Action Profile > Internal Decryption**.

4 Enter a name for the new policy.

5 (Optional) Select a profile setting: Balanced (default), High Fidelity, High Security, or use Custom to change the sub-settings.

Profile Setting	Description
Decryption Failure: Bypass & Log or Block & Log	Sets what to do when there is decryption failure which could be due to mTLS (mutual TLS) or certificate pinning in use. If you select Bypass & Log, then NSX caches this domain, and all subsequent connections to the domain are bypassed.
Crypto Enforcement: Transparent or Enforce	Sets the minimum and maximum TLS versions and cipher suites for the client and server. You can bypass this using the Transparent option.

6 (Optional) Modify Idle connection timeout. This is the time in seconds the server can remain idle after establishing a TCP connection. Default is 5400 seconds. Keep this timeout lower than the gateway firewall idle timeout settings.

7 Expand the **Server Certificates and Keys** section and configure one or more internal server certificates.

   a Import a certificate or CA certificate. See Import a Self-signed or CA-signed Certificate.

   b Generate a self-signed certificate or a self-signed CA certificate.

   c Select an existing server certificate by clicking the check box at the front of the row.

   d Make an existing server certificate the default by clicking on the **Default** radio button at the end of the row.

   If the SNI extension is not present in the client hello, the default server certificate is presented to the client. If this option is not configured, then the TLS Proxy does not intercept connections that contain no SNI extension in client hello. If the SNI extension is present in client hello, but there is no matching certificate for that in the list of configured certificates, TLS Proxy does not intercept these connections.

   When a client accesses one of those internal services, the TLS Proxy presents this selected certificate based on the server domain match to Issued-to-field (CN) field.

If more than one server certificate is configured, they must all have different domains, specified by Common Name (CN) or Subject Alternate Name (SAN). Two certificates cannot be configured for the same domain, either FQDN (for example, www.vmware.com) or wildcard (*.vmware.com). However, certificates with wildcard domains that overlap with specific FQDN certificates are allowed. In such cases, more specific certificates are preferred over wildcard certificates while selecting a certificate to present to the client based on the SNI extension in client hello, .

8 (Optional) By default the server certificate validation is optional and disabled by default. You do not need to configure this if it is an Enterprise-owned service. If you want to enforce this validation, turn on the Server Certificate Validation by toggling it **On**.

a Expand the section and configure your validation options. You can choose the default trusted CA bundle and CRL or import them.

For details, see Import or Update a Trusted CA Bundle andImport a Certificate Revocation List.

b Click **Save**.

Results

You are now able to use the internal decryption action profile to configure TLS Inspection policies and rules on your tier-1 gateways.

What to do next

Create TLS Inspection internal decryption policies and rules.

Create a Bypass Decryption Action Profile

This topic provides details about the bypass decryption action profile.

Prerequisites

Your local government and Enterprise privacy policies might forbid decryption of certain content. For example, when the client is accessing a financial website or a healthcare provider website, there might be laws forbidding interception and decryption of such traffic.

For ease of configuration, NSX includes a pre-defined context profile, default-bypass-highfidelity-profile, to meet such requirements. NSX uses context profiles to match domain URLs to be skipped, or bypassed, from decryption. The default profile includes the URL categories: healthcare and financial.

In this release, you cannot create bypass decryption action profiles or modify the default profile. The default profile has the following profile settings:

Profile Setting	Description
Invalid Certificates: Allow	Set to Allow - If the server presents with an expired or untrusted certificate, this choice allows the connection to proceed.
Crypto Enforcement: Transparent	Set to transparent - no cipher or TLS version enforcement occurs if the URL matches the bypass decryption profile rule.

TLS Inspection Certificate Management

A security certificate is required for TLS Inspection on NSX. To intercept, decrypt, and encrypt traffic for TLS Inspection and other advanced security applications, you must prepare the TLS proxy so it can act as a transparent proxy for TLS connections. NSX Manager requires these certificates to establish trust between applications.

Certificate management supports the following options for TLS Inspection.

- Import an existing certificate or generate a new self-signed or CA self-signed CSR (certificate signing request).
- Export an existing certificate.
- Import or update a default trusted CA bundle.
- Import or update a default public certificate revocation list (CRL).
- Advanced filtering on all predefined filter options.
- Expired certificate banner on Certificates page.
- Color-coded notification of valid or invalid certificates.

For information on these options, see Chapter 23 Certificates.

The TLS proxy requires a CA certificate, also called the proxy CA. NSX Manager uses the proxy CA to generate certificates that impersonate the endpoints in the intercepted connection. In other words, it helps to spoof certificates for the intercepted traffic on website servers. You can choose one of two types of proxy CA certificates:

To generate certificates that impersonate the endpoints in the intercepted connection, the TLS proxy requires a CA certificate, also called the proxy CA. NSX Manager uses the proxy CA. In other words, it helps to spoof. You can choose one of two types of proxy CA certificates:

- Self-signed certificates are typically used for testing or limited deployments that are untrusted. This workflow begins with a request to the NSX Manager to generate a CA certificate keypair with a CSR (Certificate Signing Request). It then makes a request to the NSX Manager to self-sign the CSR.

- An Enterprise issuing CA signs trusted subordinate CA certificates. This workflow begins with a request to the NSX Manager to generate a CA certificate keypair with a CSR, download the CSR, then submit the CSR to the Issuing CA which results in receiving a signed certificate. It then uploads the signed public CA certificate to the NSX Manager. The upload can include a certificate chain. Certificate chains are intermediate signing certificates between the new certificate and the root CA certificate.

There are several ways to upload a new CA certificate to the NSX Manager: use the TLS wizard in the UI, manually add the certificate in the UI, or use the NSX API. For details, see Import a CA Certificate.

Trusted CA Bundles

After setup, these CA certificates get distributed to the nodes running the TLS proxy. You can upload multiple CA certificate bundles with different names. Each CA bundle used by the TLS proxy gets configured in the decryption action profile. Upon upload, the CA bundle gets validated as a well-formed concatenation of PEM-encoded certificates and does not get stored if it is invalid. If a bundle is invalid, it returns API error messages.

A single bundle is limited to 1 MB in size and 1,000 certificates.

Certificate Revocation List

To ensure the certificates offered by the endpoints of the intercepted connection do not get revoked, you can use the TLS Proxy CRL, default_public_crl. You can update this object by uploading a new CRL to replace the existing one. You can use it in policy profiles. To upload a new CRL to the NSX Manager use the UI or API. The CRL gets distributed to the nodes running the TLS Proxy. The API validates the CRL upon upload and refuses to store it if it is invalid. NSX supports two CRL formats:

- PEM-encoded X.509 CRL - 40 MB maximum size, 500,000 entries
- Mozilla OneCRL - 5 MB maximum size, 10,000 entries

Alarm Handling for TLS Inspection Certificates

If you do not maintain your proxy CA certificates and they are close to expiring or have expired or you have received an expired CA certificate, NSX Manager uses alarms to notify you.

NSX raises the same set of alarms for expiring or expired certificates in CA bundles.

You may also receive a remote logging server error due to an invalid TLS certificate. The event error logged is `Log messages to logging server {hostname_or_ip_address_with_port} ({entity_id}) cannot be delivered possibly due to an unresolvable FQDN, an invalid TLS certificate or missing NSX appliance iptables rule.` To verify the specified certificate is valid, use the openssl command `openssl x509 -in <cert-file-path> -noout -dates`. You can also view and update certificates in the TLS Inspection UI.

For more details on certificate expiration, see Alarm Notification for Certificate Expiration. For details on TLS-specific "Certificate Events" from the NSX Manager, see the Event Catalog.

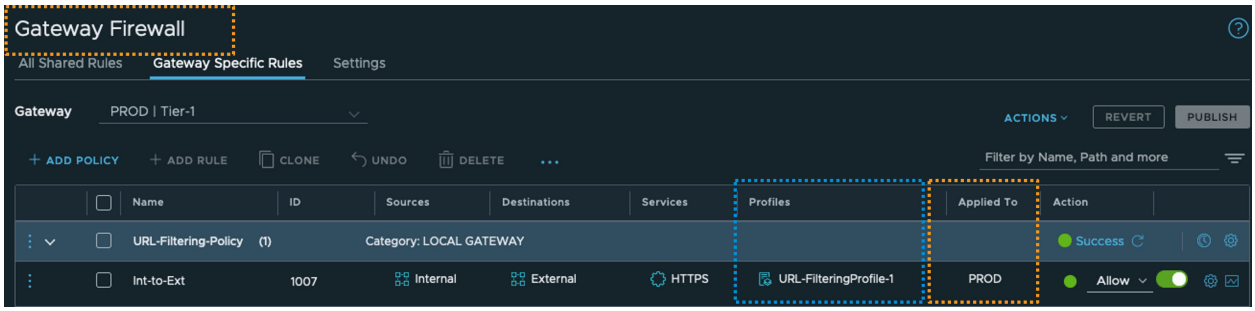
URL Filtering

URL filtering can prevent malicious code, spyware, phishing attempts and other threats by blocking access to websites or URLs that may cause a security risk. URL filtering enables access-control based on URL categories, URL reputation, and custom URLs.

URL filtering is supported only on a Tier-1 gateway. URL filtering access control policy is enforced for both encrypted (https) and non-encrypted (http) traffic. With encrypted traffic, users must enable TLS Inspection policy to decrypt the traffic to get full URL filtering. Without TLS inspection, a URL filtering policy will enforce policy at the domain level using TLS Server Name Indication (SNI) extension in the TLS client hello.

Create a URL Filtering Policy

URL filtering is configured using an L7 access profile in Gateway firewall rules.



Steps to create a URL filtering policy:

Step	Action	Refer
1.	Enable URL Database for the gateway.	Enable the URL database by navigating to **Security > General Security Settings**, and selecting the URL Database tab.
2.	Create an L7 access profile for URL filtering.	See L7 Access Profiles
3.	Create URL filtering policy	Add a Gateway Firewall Policy and Rule Create a gateway firewall rule using the URL filtering profile created in step 2. The filtering policy can only be **Applied To** Tier-1 gateways. The rule action with an L7 access profile must be **Allow.**
4. (optional)	Create a custom response or accept the default for when the L7 access profile **Action** is specified as **Reject With Response**.	See **URL Filtering** in Gateway Firewall Settings.

URL Filtering Monitoring Dashboard

The URL monitoring dashboard displays the reputation score of URLs, and other details such as the gateway, category, session count per action, and edge node.

The URL monitoring dashboard can be found by navigating to **Security > Filtering and Analysis > URL Filtering**. This dashboard does not show the URLs hitting the default entry, and App ID entries in the Layer 7 access profile used for URL filtering.

FQDN Analysis

FQDN Analysis allows administrators to gain insight into the type of websites accessed within the organization, and understand the reputation and risk of the accessed websites.

Functionality	Distributed Firewall	Gateway Firewall
FQDN Analysis (formerly URL Analysis)	NO	YES (in 3.2) ■ Supported only on tier-1 gateways. ■ Uses DNS snooping to get domains being accessed ■ Needs GFW L7 DNS rule for DNS snooping ■ DNS server needs to be north of Tier-1 gateway
URL Filtering (new in 3.2)	NO	YES ■ Uses L7 Access Profile (subset of URL filtering) ■ Uses L7 DPI/URL database for enforcement ■ HTTPS traffic: Need TLS Decryption to enforce filtering URL path
FQDN Filtering	YES ■ Using context profiles ■ Uses DNS snooping to get FQDN to IP mapping for enforcement. ■ Need DFW L7 DNS rule for DNS snooping	YES (in 3.2) ■ Uses L7 Access Profile (subset of URL filtering) ■ Uses L7 DPI/URL database for enforcement

FQDN Analysis Dashboard

The FQDN Analysis dashboard shows a summary of all analyzed URLs, classified by reputation score and category. Additionally, it provides URL specific details such as category, reputation score, gateway, edge node, and session count.

The FQDN Analysis dashboard can be found by navigating to **Security > Filtering and Analysis > FQDN Analysis**. There are more than 80 pre-defined URL categories. A site or domain can belong to multiple categories. For example, www.vmware.com belongs to both the **Business and Economy** category, and the **Computer and Internet Info** category. It is not possible to automatically drop or allow traffic based on FQDN Analysis.

Only flows traversing across a tier-1 gateway are analyzed. FQDN Analysis classifies websites into categories, and assigns a reputation score based on their domain. Based on their reputation score, URLs are classified into the following severities:

Severity Level	Description
High Risks (1-20) are red	Sites with a high probability of containing malicious links or payloads.
Suspicious (21-40) are orange	Sites with a higher than average probability of containing malicious links or payloads.
Moderate Risks (41-60) are yellow	Generally benign sites that exhibit some characteristics that suggest security risks.
Low Risks (61-80) are gray	Benign sites, that rarely exhibit characteristics that expose the user to security risks.
Trustworthy Sites (81-100) are green	Well-known sites with strong security practices.

Configuring FQDN Analysis

FQDN Analysis gives you visibilty into external domains and enables insights into cloud application usage, business elevant usage, risky user usage, and potentially malicious behavior.

Prerequisites and limitations:

- NSX Edges (management interface) need access to the internet to download category and reputation definitions from VMware Cloud.
- Medium and larger VM form factor edge nodes, or physical edge nodes are supported.
- DNS Server must be north of tier-1 gateway.
- Only north/south internet traffic from workloads deployed behind T1 is analyzed.
- Create a Layer 7 DNS rule on the tier-1 gateway to intercept DNS request and response traffic (if it doesn't already exist).

1 Navigate to **Security > Gateway Firewall** and check that you are on the **Gateway Specific** tab.

2 Click **Add Policy** to create a policy section, and give the section a name.

3 Select the check box next to the policy and click **Add Rule**.

4 Configure the following options:

Variable	Description
Name	Provide a name for the rule, such as L7 DNS Rule
Source	Any or specific group
Destination	Any or specific group
Services	■ DNS-UDP ■ DNS
Profiles	DNS context profile

Variable	Description
Applied To	Select all of the tier-1 gateways backed by the NSX Edge cluster where FQDN Analysis is enabled.
Action	Allow

5 Click **Publish**.

Activate FQDN Analysis

1 Turn on the FQDN Analysis per gateway and URL database per corresponding edge cluster by navigating to **Security > Gateway Firewall > Settings > FQDN Analysis**. Once activated, the URL database will be downloaded to each cluster member. See Gateway Firewall Settings.

Note Fetching the URL database version is not supported if a proxy server is activated in your environment. NSX Edges must have a direct internet connection with with VMware NTICS cloud to fetch the URL database version.

2 Monitor FQDN analysis on the FQDN Analysis Dashboard.

Gateway Firewall Packet Logs

If logging is enabled for a gateway firewall, gateway firewall packets will be logged.

The log file is `/var/log/firewallpkt.log`. Each log message conforms to the syslog format, and consists of a syslog header and firewall-specific information. For more informaion about syslog, see Log Messages and Error Codes.

The firewall-specifc portion of a log message has the following fields:

Field	Notes

<VRF ID and interface UUID>	You can get this information about an interface by running a CLI command. For example:
	`edge-1> get firewall interfaces`
	`Interface : 55f1af2f-4875-44e9-b0e0-59132ad7753d`
	`Type : UPLINK`
	`Sync enabled : true`
	`Name : Uplink_40_1`
	`VRF ID : 1`
	`...`

Address family	Possible values: `INET`, `INET6`

Field	Notes
Reason	Possible values: ■ `match`: Packet matches a rule. ■ `fragment`: A fragment that comes after the first fragment. ■ `short`: Packet too short (for example, no IP header, or TCP/UDP header). ■ `normalize`: Malformed packets that do not have a correct header or a payload. ■ `memory`: Datapath out of memory. ■ `ip-option`: Invalid IP options are present. ■ `TERM`: A connection is terminated.
Action	Possible values: ■ `PASS`: Accept the packet. ■ `DROP`: Drop the packet. ■ `NAT`: SNAT ■ `RDR`: DNAT ■ `PBR`: Service insertion. ■ `LB`: Load balancer.
Rule ID	The firewall rule ID.
Direction	Possible values: `IN`, `OUT`
Packet length	Length in bytes.
Protocol	Possible values: `TCP`, `UDP`, or `PROTO` (protocol number)
Source IP address and port	For SNAT, this is the address before translation.
Destination IP address and port	For DNAT, this is the address before translation.

Examples of gateway firewall log messages for TCP:

```
<181>1 2020-09-21T22:14:12.080427+00:00 lur-svc.nsxedge-ob-16404613-1-gdefw NSX 2802
FIREWALL [nsx@6876 comp="nsx-edge" subcomp="datapathd.firewallpkt" level="INFO"] <1
55f1af2f487544e9:b0e059132ad7753d> INET reason-match PASS 1005 OUT 60 TCP 1.1.1.10/45120-
>91.189.92.38/443 S

<181>1 2020-09-21T22:14:19.963758+00:00 lur-svc.nsxedge-ob-16404613-1-gdefw NSX 2802
FIREWALL [nsx@6876 comp="nsx-edge" subcomp="datapathd.firewallpkt" level="INFO"] <1
55f1af2f487544e9:b0e059132ad7753d> INET TERM PASS 1005 OUT TCP 1.1.1.10/45120-
>91.189.92.38/443
```

Examples of gateway firewall log messages for UDP:

```
<181>1 2020-09-21T22:05:05.686346+00:00 lur-svc.nsxedge-ob-16404613-1-gdefw NSX 2802
FIREWALL [nsx@6876 comp="nsx-edge" subcomp="datapathd.firewallpkt" level="INFO"] <1
55f1af2f487544e9:b0e059132ad7753d> INET reason-match PASS 1005 IN 328 UDP 40.40.40.10/60613-
>1.1.1.10/42917

<181>1 2020-09-21T22:05:48.301116+00:00 lur-svc.nsxedge-ob-16404613-1-gdefw NSX 2802
FIREWALL [nsx@6876 comp="nsx-edge" subcomp="datapathd.firewallpkt" level="INFO"] <1
55f1af2f487544e9:b0e059132ad7753d> INET TERM PASS 1005 IN UDP 40.40.40.10/60613-
>1.1.1.10/42917
```

Examples of gateway firewall log messages for PROTO:

```
<181>1 2020-09-21T21:54:38.047682+00:00 lur-svc.nsxedge-ob-16404613-1-gdefw NSX 2802
FIREWALL [nsx@6876 comp="nsx-edge" subcomp="datapathd.firewallpkt" level="INFO"] <1
55f1af2f487544e9:b0e059132ad7753d> INET reason-match PASS 1005 IN 84 PROTO 1 40.40.40.10-
>1.1.1.10

<181>1 2020-09-21T21:54:45.036957+00:00 lur-svc.nsxedge-ob-16404613-1-gdefw NSX 2802
FIREWALL [nsx@6876 comp="nsx-edge" subcomp="datapathd.firewallpkt" level="INFO"] <1
55f1af2f487544e9:b0e059132ad7753d> INET TERM PASS 1005 IN PROTO 1 40.40.40.10->1.1.1.10
```

Examples of gateway firewall log messages for SNAT:

```
<181>1 2020-09-21T22:57:24.203037+00:00 lur-svc.nsxedge-ob-16404613-1-gdefw NSX 2802
FIREWALL [nsx@6876 comp="nsx-edge" subcomp="datapathd.firewallpkt" level="INFO"] <1
55f1af2f487544e9:b0e059132ad7753d> INET reason-match PASS 1005 OUT 60 TCP 1.1.2.10/49974-
>40.40.40.10/22 S

<181>1 2020-09-21T22:57:24.203615+00:00 lur-svc.nsxedge-ob-16404613-1-gdefw NSX 2802
FIREWALL [nsx@6876 comp="nsx-edge" subcomp="datapathd.firewallpkt" level="INFO"] <1
55f1af2f487544e9:b0e059132ad7753d> INET reason-match NAT 536870914 OUT 60 TCP 2.2.2.10/37305-
OR 1.1.2.10/49974->40.40.40.10/22 S

<181>1 2020-09-21T22:57:32.125757+00:00 lur-svc.nsxedge-ob-16404613-1-gdefw NSX 2802
FIREWALL [nsx@6876 comp="nsx-edge" subcomp="datapathd.firewallpkt" level="INFO"] <1
55f1af2f487544e9:b0e059132ad7753d> INET TERM NAT 536870914 OUT TCP 2.2.2.10/37305-OR
40.40.40.10/22->1.1.2.10/49974
```

Examples of gateway firewall log messages for DNAT:

```
<181>1 2020-09-21T22:49:00.978192+00:00 lur-svc.nsxedge-ob-16404613-1-gdefw NSX 2802
FIREWALL [nsx@6876 comp="nsx-edge" subcomp="datapathd.firewallpkt" level="INFO"] <1
55f1af2f487544e9:b0e059132ad7753d> INET reason-match RDR 536870913 IN 60 TCP
40.40.40.10/40082->10.10.10.1/22-OR 1.1.1.10/22 S

<181>1 2020-09-21T22:50:01.915154+00:00 lur-svc.nsxedge-ob-16404613-1-gdefw NSX 2802
FIREWALL [nsx@6876 comp="nsx-edge" subcomp="datapathd.firewallpkt" level="INFO"] <1
55f1af2f487544e9:b0e059132ad7753d> INET TERM RDR 536870913 IN TCP 40.40.40.10/40082-
>10.10.10.1/22-OR 1.1.1.10/22
```

Distributed Security for vSphere Distributed Switch

You can install Distributed Security only for a vSphere Distributed Switch (VDS).

Installing Distributed Security for VDS provides the NSX security capabilities such as:

- Distributed Firewall (DFW)
- Distributed IDS/IPS
- Identity Firewall
- L7 App ID

- Fully Qualified Domain Name (FQDN) Filtering
- NSX Intelligence
- NSX Malware Prevention
- NSX Guest Introspection

For details about installing Distributed Security for VDS, see Install Distributed Security for vSphere Distributed Switch.

Installation Process

When you install Distributed Security, configuration changes occur only in NSX and there are no changes in VMware vCenter. The details of the VDS are discovered and the following objects are automatically created in NSX to represent the VDS details:

- A transport node profile for each cluster.
- A host switch for each VDS.
- A VLAN transport zone for each VDS.

These objects are system-generated and are not configurable or editable.

Also, as part of the VDS discovery, the Distributed Virtual port groups (DVPG) and DVports of the VDS are created as objects in NSX. For more details, see Distributed Port Groups.

Any changes made to the VDS in VMware vCenter are automatically updated in NSX.

vMotion of VMs Between Clusters With or Without Distributed Security

When you vMotion a VM from a cluster without Distributed Security to a cluster with Distributed Security, the security policies of the cluster with Distributed Security are applied to the VM.

Conversely, when you vMotion a VM from a cluster with Distributed Security to a cluster without Distributed Security, the security policies are removed.

How Upgrades to VDS Affects Distributed Security

Upgrading a VDS with Distributed Security may cause temporary disruptions to the DFW.

VDS Upgrade Path	Effect on Distributed Security
When upgrading from VDS 6.6 to any version before VDS 7.0.3.	There are no disruptions to the DFW.
When upgrading from VDS 6.6, 7.0, or 7.0.2 to VDS 7.0.3.	The DFW policies on the VDS are not enforced for a brief moment on each host during the upgrade process due to a data plane outage which occurs on the host while the upgrade is in progress. The VDS upgrade is performed concurrently across all hosts, so the overall outage period for any cluster is not significant. **Note** During the upgrade process, the DFW policies are not changed, they are only not enforced. After the upgrade process is complete on the host, the DFW policies are reinforced on the host.

Install Distributed Security for vSphere Distributed Switch

NSX allows you to install Distributed Security for vSphere Distributed Switch (VDS). As the host switch is of the type VDS, DFW capabilities can be enabled on workload VMs..

Distributed Security provides security-related functionality to your VDS such as:

- Distributed Firewall (DFW)
- Distributed IDS/IPS
- Identity Firewall
- L7 App ID
- Fully Qualified Domain Name (FQDN) Filtering
- NSX Intelligence
- NSX Malware Prevention
- NSX Guest Introspection

Prerequisites

The following are the requirements for installing Distributed Security for VDS:

- vSphere 6.7 or later.
- The vSphere cluster should have at least one VDS with distributed switch version 6.6 or later configured.
- A compute manager must be registered in NSX. See Add a Compute Manager.

Procedure

1. From your browser, log in with admin privileges to an NSX Manager at https://<nsx-manager-ip-address>.
2. Navigate to **System > Quick Start**.

3 On the **Prepare Clusters for Networking and Security** card, click **Get Started**.

4 Select the clusters that you want to install Distributed Security.

5 Click **Install NSX** and then select **Security Only**.

6 In the dialog box, click **Install**.

Note If the VDS spans across multiple clusters, Distributed Security installs only to the clusters that you selected.

The installation process for Distributed Security starts.

7 To view VDS with Distributed Security installed, navigate to **System > Fabric > Hosts**.

Note vSphere clusters prepared for Distributed Security are identified by the **Security** label.

Results

Distributed Security is installed and you can begin using security capabilities such as creating DFW policies and rules for the VDS.

Endpoint Protection

NSX allows you to insert third-party partner services as a separate service VM that provides Endpoint Protection services . A partner Service VM processes file, process, and registry events from the guest VM based on the endpoint protection policy rules applied by the NSX administrator.

Understand Endpoint Protection

Know the use case, workflow, and key concepts of endpoint protection.

Key Concepts of Endpoint Protection

The endpoint protection workflow needs partners to register their services with NSX and an administrator to consume these services. There are a few concepts that aid your understanding of the workflow.

- Service Definition: Partners define services with these attributes: name, description, supported form factors, deployment attributes that include network interfaces, and appliance OVF package location to be used by the SVM.
- Service Insertion: NSX provides the service insertion framework that allows partners to integrate networking and security solutions with the NSX platform. Guest introspection solution is one such form of service insertion.
- Service Profiles and Vendor Templates: Partners register vendor templates which expose protection levels for policies. For example, protection levels can be Gold, Silver, or Platinum. Service Profiles can be created from Vendor Templates, which allow the NSX administrators to name the Vendor Templates according to their preference. For services other than those

of Guest Introspection, the Service Profiles allow further customization using attributes. The Service Profiles can then be used in the Endpoint Protection policy rules to configure protection for virtual machine groups defined in NSX. As an administrator, you can create groups based on VM name, tags, or identifiers. Multiple Service Profiles can optionally be created from a single Vendor Template.

- Endpoint Protection Policy: A policy is a collection of rules. When you have multiple policies, arrange them in the order to run them. The same applies for rules defined within a policy. For example, policy A has three rules, and policy B has four rules, and they are arranged in a sequence such that policy A precedes policy B. When guest introspection begins running policies, rules from policy A are run first before rules from policy B.
- Endpoint Protection Rule: As an NSX administrator, you can create rules that specify the virtual machine groups that are to be protected, and choose the protection level for those groups by specifying the Service Profile for each rule.
- Service Instance: It refers to the service VM on a host. The service VMs are treated as special VMs by vCenter and they are started before any of the guest VMs are powered on and stopped after all the guest VMs are powered off. There is one service instance per service per host.

  Important Number of service instances is equal to the number of hosts on which the service is running host. For example, if you have eight hosts in a cluster, and the partner service was deployed on two clusters, the total number of service instances running are 16 SVMs.

- Service Deployment: As an admin you deploy partner Service VMs through NSX on a per cluster basis. Deployments are managed at a cluster level, so that when any host is added to the cluster, EAM automatically deploys the service VM on them.

  Automatically deploying the SVM is important because if distributed resource scheduler (DRS) service is configured on a vCenter Cluster, then vCenter can rebalance or distribute existing VMs to any new host that got added to the cluster after the SVM is deployed and started on the new host. Since partner Service VMs need NSX platform to provide security to guest VMs, the host must be prepared as a transport node.

  Important One service deployment refers to one cluster on the VMware vCenter that is managed for deploying and configuring one partner service.

- File Introspection driver: Is installed on the guest VM, intercepts the file activity on the guest VM.
- Network Introspection driver: Is installed on the guest VM, intercepts the network traffic, process, and user activity on the guest VM.

High-level Tasks for Endpoint Protection

Third-party partners services containing security scanning logic, are registered with NSX for guest VM protection. The partner service is enforced when the NSX admin deploys the registered services and applies end point protection policies to guest VM groups.

The guest introspection workflow for the endpoint protection use case is as follows:

Figure 16-7. Endpoint Protection Workflow

Workflow Tasks	Role/Persona	Implementation
Register a Service with NSX	Partner Admin	Partner Console
Register Vendor Template	Partner Admin	Partner Console
Configure Partner Services with NSX	Partner Admin	Partner Console
Deploy a Service	NSX Admin	API and NSX Manager UI
View Service Instance Details	NSX Admin	API and NSX Manager UI
Verify Health Status of Service Instance	NSX Admin	API and NSX Manager UI
Add a Service Profile	NSX Admin	API and NSX Manager UI
Consume Guest Introspection Policy	NSX Admin	API and NSX Manager UI
Add and Publish Endpoint Protection Rules	NSX Admin	API and NSX Manager UI
Monitor Endpoint Protection Status	NSX Admin	API and NSX Manager UI

Configure Endpoint Protection

Protect guest VMs running in an NSX environment using third-party partner security services.

The high-level steps to configure endpoint protection policies:

1 Ensure Prerequisites to Configure Endpoint Protection are met before you configure endpoint protection on guest VMs.

2 Supported software. See Supported Software.

3 Install File Introspection Driver for Linux VMs. See Install the Guest Introspection Thin Agent for Anti-virus on Linux Virtual Machines.

4 Install File Introspection Driver for Windows VMs. See Install the Guest Introspection Thin Agent on Windows Virtual Machines for Network Introspection.

5 Install Network Introspection Driver for Linux VMs. See Install the Linux Thin Agent for Network Introspection.

6 Create a User with Guest Introspection Partner Admin Role. See Create a User with Guest Introspection Partner Admin Role.

7 Register partner service with NSX. Refer to Partner documentation.

8 Deploy a service. See Deploy a Service.

9 Consume Guest Introspection Policy. See Consume Guest Introspection Policy.

10 Add and Publish Endpoint Protection Rules. See Add and Publish Endpoint Protection Rules.

11 Monitor endpoint protection rules. See Monitor Endpoint Protection Status.

Prerequisites to Configure Endpoint Protection

Before you configure endpoint protection for guest VMs, ensure that the prerequisites are met.

Prerequisites

- Ensure all hosts are managed by NSX Manager(s).
- Prepare and configure NSX cluster as transport nodes by applying transport node profiles. After the cluster is prepared for NSX, guest introspection components are installed. See *NSX Installation Guide*.
- Partner console is installed and configured to register services with NSX.
- Ensure that the guest VMs run VM Hardware Configuration file version 9 or higher.
- Configure VMware Tools and install thin agents.
  - See Install the Guest Introspection Thin Agent for Anti-virus on Linux Virtual Machines.
  - See Install the Guest Introspection Thin Agent on Windows Virtual Machines for Network Introspection.
  - See Install the Linux Thin Agent for Network Introspection.

Create a User with Guest Introspection Partner Admin Role

Assign a user with the Guest Introspection Partner Admin role that is available in NSX.

Note: It is recommended to register partner services by a user that is associated with the Guest Introspection Partner Admin role to avoid any security issues.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Select **System** → **User**→ **Role Assignments**.
3. Click **Add**.
4. Select the user and assign that user the **`GI Partner Admin`** role.

What to do next

Register services with NSX. See Register a Service with NSX.

Register a Service with NSX

Register third-party security services with NSX.

Prerequisites

- Ensure that prerequisites are met. See Prerequisites to Configure Endpoint Protection.
- Ensure that a vIDM user is assigned the GI Partner Admin role. This role is used to register services with NSX.

Procedure

1 Log in with the GI Partner Admin privileges to the partner console.

2 Register a service, vendor template, and configure the partner solution with NSX. See partner documentation.

What to do next

View catalog of partner services. See View Catalog of Partner Services.

View Catalog of Partner Services

The catalog page displays all the partners and their services that are registered with NSX.

Prerequisites

- Partners register services with NSX.
- Services are deployed on a cluster.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **System > Service Deployments > Catalog**.

3 Click **View** on a service. The Deployment page displays the details about the service, such as status of deployment, network details, cluster details, and so on.

What to do next

Upgrade a partner service VM.

Deploy a Service

After you register a service, you must deploy an instance of the service for the service to start processing network traffic.

Deploy partner service VMs that run the partner security engine on all the NSX hosts in a cluster. The vSphere ESX Agency Manager (EAM) service is used to deploy the partner service VMs on each host. After you deploy the SVMs, you can create policy rules used by SVM to protect guest VMs.

Prerequisites

- All hosts are managed by a VMware vCenter.
- Partner services are registered with NSX and are ready for deployment.
- NSX administrators can access partner services and vendor templates.
- Both the service VM and the partner Service Manager (console) must be able to communicate with each other at the management network level.

- Prepare cluster for NSX networking:
  - Create a transport zone.
  - Create an IP pool for tunnel endpoint IP addresses.
  - Create an uplink profile.
  - Apply transport node profile on a cluster to auto-deploy NSX on each host of the cluster.
- Starting with NSX 3.1, on clusters that span physical servers placed in different racks, you can override the transport node profile applied on a per-host basis.
- Starting with NSX 3.0, before you deploy endpoint protection service on hosts, prepare clusters by applying transport node profile.
- With NSX 2.5.x or earlier, you only need to apply transport node profile on a cluster deployed using the host-based deployment method.
- When upgrading endpoint protection service, the existing service will continue to be functional even if transport node profile is not applied to the cluster.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Go to the **System** tab and click **Service Deployment**.

3 From the Partner Service drop-down, select the service to be deployed.

4 Click **Deployment** and click **Deploy Service**.

5 Enter the service deployment name.

6 In the Compute Manager field, select the compute resource on the VMware vCenter to deploy the service.

7 In the Cluster field, select the cluster where the services need to be deployed.

8 In the Data Store drop-down menu, you can:

   a Select a datastore as the repository for the service virtual machine.

   b Select **Specified on Host**. This setting means that you do not need to select a datastore and port group on this wizard. You can directly configure agent settings on EAM in VMware vCenter to point to a specific datastore and port group to be used for service deployment.

   To know how to configure EAM, refer to the vSphere documentation.

9 In the Network column, click **Set**.

10 Set the Management Network interface to **Specified on Host** or **DVPG**.

11 Set the network type to DHCP or Static IP pool. If you set the network type to Static IP pool, select from the list of available IP pools.

12 In the Deployment Specification field, select form factor of the service for deployment on all hosts.

13 In the Deployment Template field, select the registered deployment template.

14 Click **Save**.

Results

When a new host is added to the cluster, EAM automatically deploys the service VM on the new host. The deployment process might take some time, depending on the vendor's implementation. You can view the status in the NSX Manager user interface. The service is successfully deployed on the host when the status turns `Deployment Successful`.

To remove host from a cluster, first move it into maintenance mode. Then, select the option to migrate the guest VMs to another host to complete migration.

What to do next

Know deployment details and heath status about service instances deployed on hosts. See View Service Instance Details.

View Service Instance Details

Know deployment details and health status of service instance deployed on member hosts of a cluster.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **System > Service Deployments > Service Instances**.

3 From the Partner Service drop-down menu, select the partner service to view details related to service instances.

Table 16-5.

Field	Description
Service Instance Name	A unique ID identifying the service instance on a particular host.
Service Deployment Name	The name you entered when deploying the service.
Deployed To	Host IP address or FQDN
Deployment Mode	Cluster or Standalone

Table 16-5. (continued)

Field	Description
Deployment Status	Up status to determine a successful deployment
Health Status	When the service instance is deployed, the heath status is `Ready`. To bring the health status from `Ready` to `Up`, make the required configuration changes. See Verify Health Status of Service Instance. After the following parameters are successfully realized by NSX, the health status changes from `Ready` to `Up`. ■ Solution status: Up ■ Connectivity between NSX Guest Introspection agent and NSX Ops Agent: Up ■ Health Status received at: <Day, Date, Time>

What to do next

Bring up Service Instance. See Verify Health Status of Service Instance.

Verify Health Status of Service Instance

After deploying the service instance, certain parameters need to be realized in NSX for the health status to be Up.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **System > Service Deployments > Service Instances**.

3 From the Partner Service drop-down menu, select the partner service to view details related to service instances.

4 The Health Status column displays state of the service instance as `Ready`. It indicates that the service instance is ready to be configured with endpoint protection policy rules to protect VMs.

5 The following parameters must be realized in NSX for the health status to change to `Up`.

- Guest virtual machines must be available on the host.
- Guest virtual machines must be powered on.
- Endpoint protection rules must be applied to the guest virtual machines.
- Guest virtual machines must be configured with the supported version of VMtools and file introspection drivers.

What to do next

Add a service profile. See Add a Service Profile.

Add a Service Profile

Guest introspection policies can be implemented only when a service profile is available in NSX. Service profiles are created from a template provided by the partner. Service Profiles are a way for the administrator to choose protection levels (Gold, Silver, Platinum policy) for a VM by choosing the vendor templates provided by the vendor.

For example, a vendor can provide Gold, Platinum, and Silver policy levels. Each profile created might serve a different type of workload. A Gold service profile provides complete antimalware to a PCI-type workload, while a silver service profile only provides basic antimalware protection to a regular workload.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Select **Security > Endpoint Protection Rules > Service Profiles** .
3. From the Partner Service field, select the service for which you want to create a service profile.
4. Click **Add Service Profile**.
5. Enter the service profile name and select the vendor template. Optionally, add description and tags.
6. Click **Save**.

   The vendor template ID used to create the service profile is passed on to the partner console. Partners store the vendor template ID to track usage of which guest VMs are protected by these vendor template.

Results

After creating service profile, an NSX admin creates rules to associate a service profile to a group of VMs before publishing the policy rule.

What to do next

Apply endpoint protection policy on guest VM groups that need to be protected from malware. See Consume Guest Introspection Policy.

Consume Guest Introspection Policy

Policy can be enforced on VM groups by creating rules that associate service profiles with VM groups. Protection begins immediately after rules are applied to a VM group.

The endpoint protection policy is a protection service offered by partners to protect guest VMs from malware by implementing service profiles on guest VMs. With a rule applied to a VM group, all guest VMs within that group are protected by that service profile. When a file access event on a guest VM occurs, the GI thin agent (running on each guest VM) collects context of the file (file

attributes, file handle, and other context details) and notifies the event to SVM. If the SVM wants to scan the file content, it request for details using the EPSec API library. Upon a clean verdict from SVM, the GI thin agent allows the user to access the file. In case SVM reports the file as infected, the GI thin agent denies user access to the file.

To execute an security service on a VM group, you need to:

Procedure

1 Define policy and rules.

2 Define membership criteria to form VM group.

3 Define rules for VM groups.

4 Publish the rule.

Add and Publish Endpoint Protection Rules

Publishing endpoint protection policy rules to VM groups means associating VM groups that need to be protected with a specific service profile.

Endpoint protection policy rules can only be applied to VM groups.

Procedure

1 Select **Security > Endpoint Protection Rules > Rules > Add Policy**.

2 Click  .

3 Select **Add Rule**.

4 In the new rule, enter the rule name.

5 In the Select Groups field, click the Edit icon.

6 In the Set Groups window, select from the existing list of groups or add a new group.

   a To add a new group, click **Add Group**, enter details and click **Save**.

   See Add a Group.

7 In the Group column, select the VM group.

8 In the Service Profiles column, select the service profile that provides the desired protection level to the guest VMs in the group.

   a To add a new service profile, click **Add Service Profile**, enter details and click **Save**.

   See Add a Service Profile.

9 Click **Publish**.

Results

Endpoint protection policies protect VM groups.

What to do next

You might want to change the sequence of rules depending on the type of protection required for different VM groups. See How Guest Introspection Runs Endpoint Protection Policy

Monitor Endpoint Protection Status

Monitor the configuration status of protected and unprotected VMs, issues with Host agent and service VMs, and VMs configured with the file introspection driver that was installed as part of the VMtools installation.

You can view:

- View Service Deployment Status.
- View Configuration Status of Endpoint Protection.
- View Capacity Status Set for Endpoint Protection.

View Service Deployment Status

View service deployment details on the Monitoring Dashboard.

View the system-wide status of EPP policy.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Navigate to **Home** > **Monitoring - Dashboards**.
3. From the drop-down menu, click **Monitoring - System**.
4. To view the deployment status across clusters in the system, navigate to the Endpoint Protection widget, click the doughnut chart to view successful or unsuccessful deployments.

   The Service Deployments page displays the deployment details.

View Configuration Status of Endpoint Protection

View configuration status of the endpoint protection service.

View the system-wide status of EPP policy.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Navigate to **Home** > **Security** > **Security Overview**.
3. To view status of EPP on clusters, click the Security widget.
4. In the Security Overview page, click **Configuration**.



5 In the Endpoint Protection section, view:

a VM Distribution by Service Profile widget displays:

1 Number of VMs protected by top profile. Top profile represents a profile that protects the maximum number of VMs on a cluster.

2 VMs protected by remaining service profiles categorized under Other Profiles.

3 VMs not protected categorized under No Profile.

The Endpoint Protection Rules page displays VMs protected by Endpoint Protection policies.

b Components having issues widget displays:

1 Host: Issues related to the context multiplexer.

2 SVM: Issues related to service VMs. For example, the SVM state is down, SVM connection with guest VM is down.

The Status column on the Deployment page displays health issues.

c Configure VMs running File Introspection widget displays:

1 VMs protected by File Introspection driver.

2 VMs where the File Introspection driver status is unknown.

ESXi Agency Manager (EAM) attempts to resolve a few issues related to hosts, SVMs, and configuration errors. See Resolve Partner Services Issues.

View Capacity Status Set for Endpoint Protection

View capacity status of the endpoint protection service.

View the capacity status of EPP policy.

Procedure

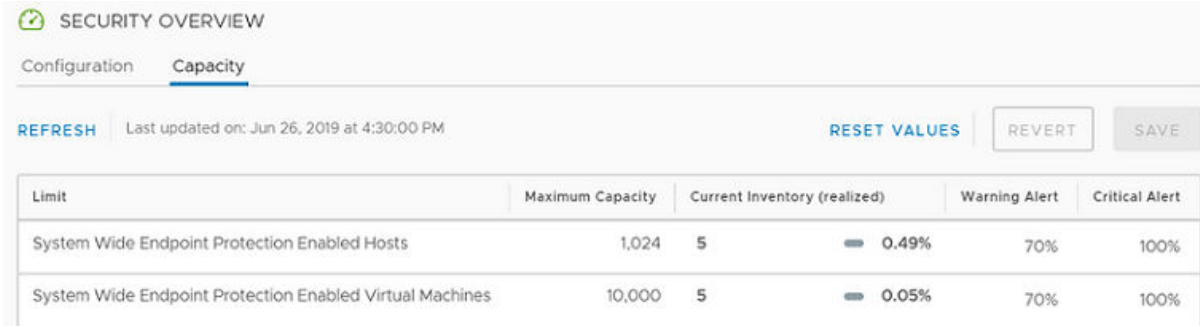
1 With admin privileges, log in to NSX Manager.

2 Navigate to **Home** > **Monitoring - Dashboards**.

3 From the drop-down menu, click **Monitoring - Networking and Security**.

4 To view status of EPP on clusters, click the Security widget.

5 In the Security Overview page, click **Capacity** and view capacity status of these parameters.



a **System Wide Endpoint Protection Enabled Hosts**: If the number of host numbers protected reaches the threshold limit, NSX Manager notifies a warning alert or critical alert when corresponding threshold limits are reached.

b **System Wide Endpoint Protection Enabled Virtual Machines**: If the number of virtual machine numbers protected reaches the threshold limit, NSX Manager notifies a warning alert or critical alert when corresponding threshold limits are reached.

Note You can set threshold limits for these parameters, view status and receive alerts when these parameters reach the set threshold limit.

Change the Third Party Service Virtual Machine

An NSX administrator can change or deploy a new form factor or version of the service VM (SVM).

Note In releases before NSX 3.2, the **Change Appliance** was called **Re-deploy Appliance**.

This task can be done both from UI or API.

The API command to change or upgrade a SVM is `/ POST https://<NSX_Manager_IPaddress>/api/v1/serviceinsertion/services/{{service_id}}/ service-deployments /<service-deployment-id>?action=upgrade`.

```
{
   "deployment_spec_name": "EW_DepSpec"
}
```

Prerequisites

- Ensure that partners have registered multiple service VMs differentiated by versions and or form factor (disk, vCPU, or RAM). The **Change Appliance** is available only when multiple service VMs are deployed.
- Ensure that the SVM deployment status is `Deployment Successful` before changing the appliance. If the SVM is in a different state, go to **Home** → **Alarms**, and search for any open alarms of the event type EAM. Resolve errors before trying to change to a newer SVM.

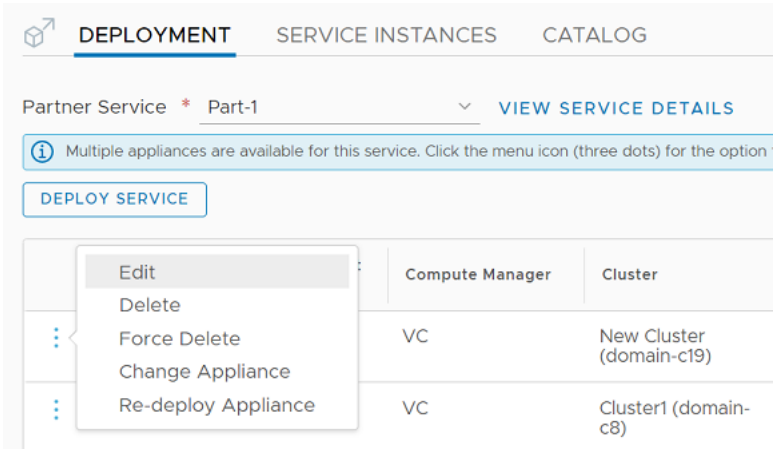
- Ensure that all prereqiuistes required to deploy an endpoint service or a combined partner service (for example, endpoint protection service and network gateway firewall) are met before proceeding to change the appliance.
- Ensure that storage is available before you change the existing SVM with a new one.
- If there are workloads that are protected by the existing SVM, first perform vMotion to migrate the workload and then change or deploy the new SVM.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Go to **System > Service Deployments > Deployment**.

3 Go to the service deployment and click **Change Appliance**.



4 In the **Change Appliance** window, select the specification of SVM you want to deploy and click **Update**.

5 If the service deployed fails at an individual host-level, in the Status column, click the information icon and click **Resolve**.

6 Click **Resolve All**.

7 If the service deployed fails at cluster-level, click the vertical ellipsis and click **Re-deploy Appliance**.

Results

What to do next

To know the runtime status of a combined partner service, health status of an endpoint protection service, or deployment status, go to the **Service Instances** tab.

After changing the existing SVM to a new SVM, enable guest VMs that are to be protected by the endpoint protection policy service.

Manage Endpoint Protection

Resolve policy conflicts, health issues with service VMs, and know how endpoint protection policy works.

Resolve Partner Services Issues

Without partner service virtual machine functional, guest VMs are not protected against malware.

On each host, verify that the following services or process are up and running:

- ESXi Agency Manager (EAM) service must be up and running. The following URL must be accessible.

  ```
  https://<vCenter_Server_IP_Address>/eam/mob
  ```

  Verify the ESXi Agency Manager is online.

  ```
  root> service-control --status vmware-eam
  ```

- Port groups of SVMs must not be deleted because these port groups are required to ensure that SVM continues to protect guest VMs.

  ```
  https://<vCenter_Server_IP_Address>/ui
  ```

- In VMware vCenter, go to the virtual machine, click the **Networks** tab, and check whether **vmservice-vshield-pg** is listed.
- Context Multiplexer (MUX) service is up and running. Check `nsx-context-mux` VIB is UP and running on the host.
- The management interface on which NSX communicates with the partner service console must be up.
- The control interface enabling communication between MUX and SVM must be up. Port group connecting MUX with SVM must be created. Both interface and port group are required for the partner service to be functional.

ESXi Agency Manager Issues

The table lists the ESXi Agency Manager issues that can be resolved using the Resolve button on the NSX Manager user interface. It notifies NSX Manager with error details.

Table 16-6. ESXi Agency Manager Issues

Issue	Category	Description	Resolution
Cannot Access Agent OVF	VM Not Deployed	An agent virtual machine is expected to be deployed on a host, but the agent virtual machine cannot be deployed because the ESXi Agent Manager is unable to access the OVF package for the agent. It might happen because the web server providing the OVF package is down. The web server is often internal to the solution that created the Agency.	ESXi Agency Manager (EAM) service retries the OVF download operation. Check the partner management console status. Click **Resolve**.
Incompatible Host Version	VM Not Deployed	An agent virtual machine is expected to be deployed on a host. However, because of compatibility issues the agent was not deployed on the host.	Upgrade either the host or the solution to make the agent compatible with the host. Check the compatibility of the SVM. Click **Resolve**.
Insufficient Resources	VM Not Deployed	An agent virtual machine is expected to be deployed on a host. However, ESXi Agency Manager (EAM) service did not deploy the agent virtual machine because the host has less CPU or memory resources.	ESXi Agency Manager (EAM) service attempts to redeploy the virtual machine. Ensure that CPU and memory resources are available. Check the host and free up some resources. Click **Resolve**.
Insufficient Space	VM Not Deployed	An agent virtual machine is expected to be deployed on a host. However, the agent virtual machine was not deployed because the agent datastore on the host did not have enough free space.	ESXi Agency Manager (EAM) service attempts to redeploy the virtual machine. Free up some space on the datastore. Click **Resolve**.
No Agent VM Network	VM Not Deployed	An agent virtual machine is expected to be deployed on a host, but the agent cannot be deployed because the agent network has not been configured on the host.	Add one of the networks listed in customAgentVmNetwork to the host. The issue resolves automatically after the datastore is available.

Table 16-6. ESXi Agency Manager Issues (continued)

Ovf Invalid Format	VM Not Deployed	An Agent virtual machine is expected to be provisioned on a host, but it failed to do so because the provisioning of the OVF package failed. The provisioning is unlikely to succeed until the solution that provides the OVF package has been upgraded or patched to provide a valid OVF package for the agent virtual machine.	ESXi Agency Manager (EAM) service attempts to redeploy the SVM. Check the partner solution documentation or upgrade the partner solution to get the valid OVF package. Click **Resolve**.
Missing Agent IP Pool	VM Powered Off	An agent virtual machine is expected to be powered on, but the agent virtual machine is powered off because there are no IP addresses defined on the agent's virtual machine network.	Define the IP address on the virtual machine network. Click **Resolve**.
No Agent VM Datastore	VM Powered Off	An agent virtual machine is expected to be deployed on a host, but the agent cannot be deployed because the agent datastore has not been configured on the host.	Add one of the datastores listed in customAgentVmDatastore to the host. The issue resolves automatically after the datastore is available.
No Custom Agent VM Network	No Agent VM Network	An agent virtual machine is expected to be deployed on a host, but the agent cannot be deployed because the agent network has not been configured on the host.	Add the host to one of the networks listed in a custom agent VM network. The issue resolves automatically after a custom VM network is available.
No Custom Agent VM Datastore	No Agent VM Datastore	An agent virtual machine is expected to be deployed on a host, but the agent cannot be deployed because the agent datastore has not been configured on the host.	Add the host to one of the datastores listed in a custom agent VM datastore. The issue resolves automatically.
Orphaned Agency	Agency Issue	The solution that created the agency is no longer registered with the VMware vCenter.	Register the solution with the VMware vCenter.

Table 16-6. ESXi Agency Manager Issues (continued)

Orphaned DvFilter Switch	Host Issue	A dvFilter switch exists on a host but no agents on the host depend on dvFilter. It happens if a host is disconnected when an agency configuration changed.	Click **Resolve**. ESXi Agency Manager (EAM) service attempts to connect the host before the agency configuration is updated.
Unknown Agent VM	Host Issue	An agent virtual machine has been found in theVMware vCenter inventory that does not belong to any agency in this vSphere ESX Agent Manager server instance.	Click **Resolve**. ESXi Agency Manager (EAM) service attempts to place the virtual machine to the inventory it belongs to.
Ovf Invalid Property	VM Issue	An agent virtual machine must be powered on, but an OVF property is either missing or has an invalid value.	Click **Resolve**. ESXi Agency Manager (EAM) service attempts to reconfigure the correct OVF property.
VM Corrupted	VM Issue	An agent virtual machine is corrupt.	Click **Resolve**. ESXi Agency Manager (EAM) service attempts to repair the virtual machine.
VM Orphaned	VM Issue	An agent virtual machine exists on a host, but the host is no longer part of scope for the agency. It happens if a host is disconnected when the agency configuration is changed.	Click **Resolve**. ESXi Agency Manager (EAM) service attempts to connect the host back to the agency configuration.
VM Deployed	VM Issue	An agent virtual machine is expected to be removed from a host, but the agent virtual machine has not been removed. The specific reason why vSphere ESX Agent Manager was unable to remove the agent virtual machine, such as the host is in maintenance mode, powered off or in standby mode.	Click **Resolve**. ESXi Agency Manager (EAM) service attempts to remove the agent virtual machine from the host.
VM Powered Off	VM Issue	An agent virtual machine is expected to be powered on, but the agent virtual machine is powered off.	Click **Resolve**. ESXi Agency Manager (EAM) service attempts to power on the virtual machine.

Table 16-6. ESXi Agency Manager Issues (continued)

VM Powered On	VM Issue	An agent virtual machine is expected to be powered off, but the agent virtual machine is powered off.	Click **Resolve**. ESXi Agency Manager (EAM) service attempts to power off the virtual machine.
VM Suspended	VM Issue	An agent virtual machine is expected to be powered on, but the agent virtual machine is suspended.	Click **Resolve**. ESXi Agency Manager (EAM) service attempts to power on the virtual machine.
VM Wrong Folder	VM Issue	An agent virtual machine is expected to be located in a designated agent virtual machine folder, but is found in a different folder.	Click **Resolve**. ESXi Agency Manager (EAM) service attempts to place the agent virtual machine to the designated folder.
VM Wrong Resource Pool	VM Issue	An agent virtual machine is expected to be located in a designated agent virtual machine resource pool, but is found in a different resource pool.	Click **Resolve**. ESXi Agency Manager (EAM) service attempts to place the agent virtual machine to a designated resource pool.
VM Not Deployed	Agent Issue	An agent virtual machine is expected to be deployed on a host, but the agent virtual machine has not been deployed. Specific reasons why ESXi Agent Manager was unable to deploy the agent, such as being unable to access the OVF package for the agent or a missing host configuration. This issue can also happen if the agent virtual machine is explicitly deleted from the host.	Click **Resolve** to deploy the agent virtual machine.

NSX Manager Issues

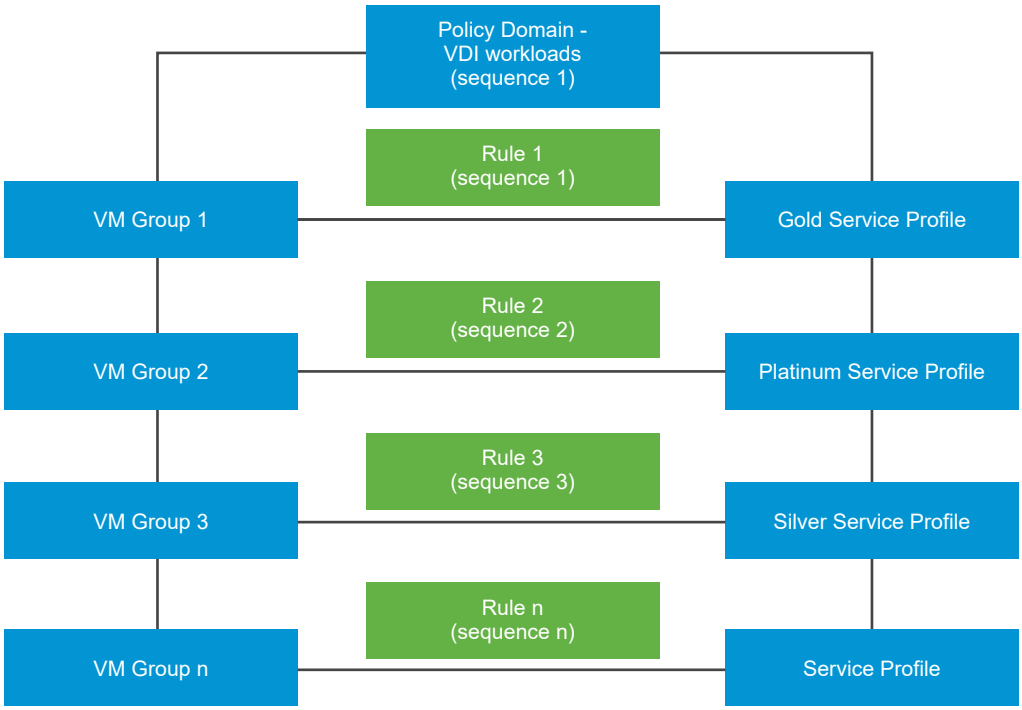
Issue	Description	Resolution
Unable to allocate static IP addresses from the IP Pool	Either the IP addresses from the pool are exhausted or there are no more IP addresses left to allocate.	Fix the IP Pool problem, click **Resolve** to fix the issue.
OVF certification error	NSX was not able to certify the OVF provided in the service. Either the certificates are not valid or the location is not reachable.	Verify whether the OVF is certified. Verify whether the OVF/Certificate file location is reachable from the NSX Manager appliance. After verifying the above points, delete the deployment and start a new deployment.

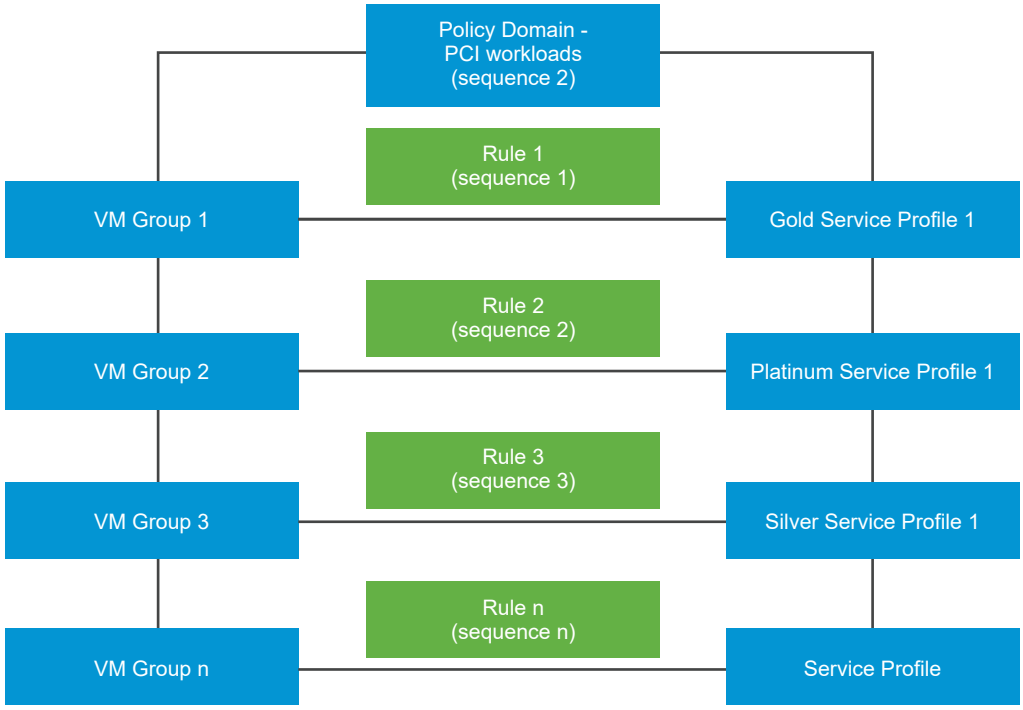
Next, configure the Endpoint Protection for VM groups. See Endpoint Protection.

How Guest Introspection Runs Endpoint Protection Policy

Endpoint protection policies are enforced in a specific order. When you design policies, consider the sequence number associated to rules and the domains that host the rules.

Scenario: Out of the many workloads that run in your organization, for the purposes of illustration we consider two kinds of workloads - VMs running Virtual Desktop Infrastructure (VDI), and VMs running Payments Cards Industry Data Security Standards (PCI-DSS) workloads. A section of employees in the organization requires remote desktop access, which makes up the virtual desktop infrastructure (VDI) workload. These VDI workloads might require a Gold protection policy level based on the compliance rules set up by the organization. Whereas a PCI-DSS workload needs the highest level of protection, Platinum level protection.





As there are two workload types, create two policies one each for VDI workloads and server workloads. Within each policy or section, define a domain to reflect the workload type and within that section define rules for that workload. Publish the rules to start GI services on guest VMs. GI internally uses the two sequence numbers: Policy sequence number and rule sequence number to determine the complete sequence of rules to run. Each rule serves two purposes: determines which VMs to protect and the protection policy that must be applied to protect the VMs.

To change the sequence order, drag a rule in the NSX Policy Manager UI to change its sequence order. Alternatively, you can explicitly assign sequence number for rules using API.

Alternatively make an NSX API call to manually define a rule by associating a service profile with a VM group and declare the sequence number of the rules. The API and parameter details are detailed in the NSX *API guide*. Make Service configuration APIs calls to apply profiles to entities such as VM groups and so on.

Table 16-7. NSX APIs used to define rule that apply service profile to VM groups

API	Details
Get all service configuration details.	`GET /api/v1/service-configs` The service configuration API returns details of the service profile applied to a VM group, the VM group protected, and the sequence or precedence number that decides priority of the rule.
Create a service configuration.	`POST /api/v1/service-configs` The service configuration API takes input parameters of a service profile, VM group to be protected, and sequence or precedence number that must be applied to the rule.
Delete a service configuration.	`DELETE /api/v1/service-configs/ <config-set-id>` The service configuration API deletes the configuration applied to the VM group.
Get details of a specific configuration.	`GET /api/v1/service-configs/ <config-set-id>` Get details of a specific configuration
Update a service configuration.	`PUT /api/v1/service-configs/ <config-set-id>` Update a service configuration.
Get effective profiles.	`GET /api/v1/service-configs/ effective-profiles?resource_id=<resource-id> &resource_type=<resource-type>` The service configuration API returns only that profile which is applied to a particular VM group.

Efficiently manage rules by following these recommendations:

- Set a higher sequence number for a policy for which rules must be ran first. From the UI, you can drag policies to change their priority.
- Similarly, set a higher sequence number for rules within each policy.
- Depending on how many rules you need, you can position rules apart in multiples of 2, 3, 4, or even 10. So, two consecutive rules that are 10 positions apart give you more flexibility to resequence rules without having to change the sequence order of all the rules. For example, if you do not plan to define many rules, you can select to position rules 10 positions apart. So, rule 1 gets a sequence number of 1, rule 2 gets a sequence number of 10, rule 3 gets a sequence number of 20, and so on. This recommendation provides flexibility to efficiently manage rules so that you do not need to resequence all the rules.

Internally, guest introspection sequences these policy rules in the following way.

```
Policy 1 ↔ Sequence Number 1 (1000)

 - Rule 1 : Group 1↔ Service Profile ↔ Sequence Number 1 (1001)

 - Rule 2 : Group 1↔ Service Profile ↔ Sequence Number 10 (1010)

 - Rule 3 : Group 1↔ Service Profile ↔ Sequence Number 20 (1020)

 - Rule 4 : Group 1↔ Service Profile ↔ Sequence Number 30 (1030)



Policy 2  ↔ Sequence Number 2 (2000)

 - Rule 1 : Group 1↔ Service Profile ↔ Sequence Number 1 (2001)

 - Rule 2 : Group 1↔ Service Profile ↔ Sequence Number 10 (2010)

 - Rule 3 : Group 1↔ Service Profile ↔ Sequence Number 20 (2020)

 - Rule 4 : Group 1↔ Service Profile ↔ Sequence Number 30 (2030)
```

Based on the above sequence numbers, GI runs rules of Policy 1 before it runs rules of Policy 2.

But there are situations when the intended rules are not applied to a VM group or a VM. These conflicts need to be resolved to apply the desired policy protection levels.

Endpoint Policy Conflict Resolution

Consider a scenario where two policy domains exist, each consisting of multiple rules. As an admin you are not always certain of which VMs can end up getting membership of a group because VMs get associated to a group based on dynamic membership criteria, such as OS Name, Computer Name, User, Tagging.

Conflicts arise in the following scenarios:

- A VM is part of two groups, where each group is protected by a different profile.
- A partner service VM is associated with more than one service profile.
- An unexpected rule ran on a guest VM, or when a rule does not run on a VM group.
- Sequence number is not assigned to policy rules or domains.

Table 16-8. Resolve policy conflicts

Scenario	Expected Endpoint Protection Flow	Resolution
When a VM gets membership to multiple groups. And each group is protected by a different type of service profile. Expected protection was not applied to the VM.	A VM group created with a membership criteria means that VMs are added to the group dynamically. In such a case, the same VM can be part of multiple groups. There is no way to pre-determine which group that VM is going to be part of because the membership criteria dynamically populates VM into the group. Consider VM 1 is part of Group 1 and Group 2. ■ Rule 1: Group 1 (by OS name) is applied Gold (Service Profile) with Sequence Number 1 ■ Rule 2: Group 2 (by tag) is applied Platinum with Sequence Number 10 Endpoint protection policy runs the Gold service profile on VM 1 but does not run Platinum service profile on VM1.	Change the Sequence Number of Rule 2 such that it runs before Rule 1. ■ On the NSX Policy Manager UI, drag the Rule 2 before Rule 1 on the rule list. ■ Using NSX Policy Manager API, manually add a higher sequence number for Rule 2.
When a rule associates the same service profile to protect two VM groups. Endpoint protection does not run the rule on the second VM group.	Endpoint protection only runs the first service profile on the VM because the same service profile cannot be applied again to any other rule across policies or domain. Consider VM 1 is part of Group 1 and Group 2. Rule 1: Group 1 (by OS name) is applied Gold (service profile) Rule 2: Group 2 (by tag) is applied Gold (service profile)	■ Add Group 2 to Rule 1. (Rule 1: Group 1, Group 2 is applied Profile 1)

Quarantine VMs

After rules are applied to VM groups, based on the protection level and tag set by partners, there might be VMs that are identified as infected that need to be quarantined.

Partners use the API with tag `virus_found=true` to tag VMs that are infected. Affected VMs are attached with the `virus_found=true` tag.

As an administrator, you can create a pre-defined quarantine group based on tag with `virus_found=true` value, such that the group gets populated with infected VMs as and when they are tagged. As an admin, you might choose to set specific firewall rules for the quarantine group. You can set firewall rules for the quarantine group. For example, you might choose to block all traffic incoming and outgoing from the quarantine group.

Verify Health Status of Service Instances

Health status of a service instance depends on many factors: status of the partner solution, connectivity between Guest Introspection Agent (Context Multiplexer) and Context Engine (Ops

Agent), and availability of Guest Introspection Agent information, SVM protocol information with NSX Manager.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **System > Service Deployments > Service Instances**.

3 In the Health Status column, click  to know the health of the service instance.

Table 16-9. Health Status of Third-Party Service Instance

Parameter	Description
Health Status received at	The latest timestamp when NSX Manager received the health status details of the service instance.
Solution Status	Status of partner solution running on an SVM. Status UP indicates that the partner solution is correctly running.
Connectivity between NSX Guest Introspection Agent and NSX Ops Agent	Status is UP when NSX Guest Introspection agent (context multiplexer) is connected with the Ops agent (includes the context engine). The context multiplexer forwards health information of SVMs to the context engine. They also share SVM-VM configuration between each other to know which guest VMs are protected by the SVM.
Service VM Protocol Version	Transport protocol version used internally for troubleshooting issues.
NSX Guest Introspection Agent Information	Represents protocol version compatibility between NSX Guest Introspection agent and SVM.

4 If the Health Status is `Up` (status displayed in green) and the partner console displays all guest VMs as protected, the health status of the service instance is `Up`.

5 If the Health Status is `Up` (status displayed in green) but the partner console displays guest VMs in unprotected state, perform the following step:

a Contact VMware support to resolve the issue. The health status of the service instance might be down not correctly reflected by the NSX Manager user interface.

6 If the Health Status is `Down` (status displayed in red), then one or more factors that determine the service instance health are down.

Table 16-10. Troubleshoot Health Status

Health Status Attribute	Resolution
Solution Status is `Down` or `Not available`.	1 Verify that service deployment status is `Up` (green). If you encounter errors, see Resolve Partner Services Issues. 2 Ensure that at least one guest VM in the affected host is protected with an endpoint protection policy. 3 From the partner console, verify whether the solution service is running on the SVM on the host. See the Partner documentation for more details. 4 If none of the above steps resolve the issue, contact VMware support.
Connectivity between NSX Guest Introspection Agent and NSX Ops Agent is `Down`.	1 Verify that service deployment status is `Up` (green). If you encounter errors, see Resolve Partner Services Issues. 2 Ensure that at least one guest VM in the affected host is protected with an endpoint protection policy. 3 From the partner console, verify whether the solution service is running on the SVM on the host. See the Partner documentation for more details. 4 If none of the above steps resolve the issue, contact VMware support.
Service VM Protocol Version is `Unavailable`.	1 Verify that service deployment status is `Up` (green). If you encounter errors, see Resolve Partner Services Issues. 2 Ensure that at least one guest VM in the affected host is protected with an endpoint protection policy. 3 From the partner console, verify whether the solution service is running on the SVM on the host. See the Partner documentation for more details. 4 If none of the above steps resolve the issue, contact VMware support.
NSX Guest Introspection Agent Information is `Unavailable`.	Contact VMware support.

Delete Partner Services

Delete partner services through NSX Manager UI or API call.

Before you delete partner services or SVMs deployed on a host, you need to do the following from the NSX Manager UI.

To delete partner services:

Procedure

1 Remove EPP rules applied to VM groups running on the host.

2 Remove service profile protection applied to VM groups.

3 Navigate to **System > Service Deployments > Deployment**.

4 From the **Partner Service** drop-down menu, select the partner service.

5 Click the vertical ellipses icon of the service you want to delete.

6 Click **Delete**. The service will be permanently deleted and cannot be recovered.

7 In the pop-up window, click **Delete**.

8 If the service cannot be deleted, the Status goes into `Failed` state.

9 To completely delete a service that is in `Failed` state, click the vertical ellipses, and click **Force Delete**.

10 Alternatively, to remove solution binding SVMs with partner service manager, execute the following API call.

```
/DEL https://<NSX_Manager_IPaddress>/api/v1/serviceinsertion/services/{{service_id}}/
solution-configs/<solution-config-id>
```

11 To delete the service deployment, execute the following API call.

```
/DEL https://<NSX_Manager_IPaddress>/api/v1/serviceinsertion/services/<service-id>/service-
deployments/<service-deployment-id>
```

Refer to the *NSX API Guide* guide for more information on API parameters.

Troubleshooting Endpoint Protection

Table 16-11. Endpoint Protection Troubleshooting

Symptoms	Resolution and Troubleshooting Steps
Clusters not ready after trying to change the IP Pool for a cluster.	Register the certificate with vCenter.
Host preparation displayed as "not ready."	EAM was down due to a buffer overflow - increase buffer overflow to 256 and reboot.
Unable to install NSX VIBs.	Update EAM Solution's Certificate (KB 2112577) to deploy EAM Agency. Add vCenter's Managed IP Address (KB 1008030).
Can't prepare an NSX Cluster.	Restart EAM agent in vCenter.

Collecting Endpoint Protection Environment and Work Details

Collecting environment details is useful when checking the compatibility of components.

1 Determine if Endpoint Protection is used in the customer environment. If it is not, remove the Endpoint Protection service for the virtual machine, and confirm the issue is resolved.

2 Collect environment details:

   a ESXi build version - Run the command `uname –a` on the ESXi host or click on a host in the vSphere Web Client and look for the build number at top of the right-hand pane.

   b Linux product version and build number

   c `/usr/sbin/vsep –v` will give the production version

```
Build number
------------------
Ubuntu
dpkg -l | grep vmware-nsx-gi-file
SLES12 and RHEL7
rpm -qa | grep vmware-nsx-gi-file
```

3 Collect NSX version, and the following details:

   - Partner solution name and version number
   - EPSec Library version number used by the partner solution: Log into the EPP SVM and run this command:

```
strings <path to EPSec library>/libEPSec.so | grep BUILD
```

   - Guest operating system in the virtual machine
   - Any other third-party applications or file system drivers

4 ESX EPP Module (MUX) version - run the command esxcli software `vib list | grep nsx-context-mux`.

5 Collect workload details, such as the type of server.

6 Collect ESXi host logs. For more information, see Collecting diagnostic information for VMware ESX/ESXi (653).

7 Collect service virtual machine (EPP SVM) logs from the partner solution. Reach out to your partner for more details on EPP SVM log collection.

8 Collect a suspend state file while the problem is occurring, see Suspending a virtual machine on ESX/ESX (2005831) to collect diagnostic information.

9 After collecting date, compare the compatibility of the vSphere components. For more information, see the VMware Product Interoperability Matrices.

Service Status Unknown or Endpoint Protection fails to get IP Address

After deployment, Endpoint Protection (EPP) service has an “unknown” status in vCenter or the Endpoint Protection VM does not receive an IP address.

Problem

After deployment, Endpoint Protection status shows as “Not Ready." EPP is otherwise reachable with valid IP assigned. “Trend SVM Heartbeat status” shows as red.

Cause

Deployed Endpoint Protection does not have a valid IP address.

Endpoint Protection Service Deployment shows status as “Unknown”

Solution

1. If Endpoint Protection is lacking an IP address or shows as Failed: Networking:
   a. On the host, ensure that each host has been configured properly. See the *NSX Installation Guide*.
   b. Deploy vSwitch and distributed port group for EPP. EPP should be deployed on a DVPortGroup created for the network on an existing NSX vSwitch.
   c. Ensure that the physical firewall and existing network configuration is valid.
2. If Endpoint Protection is lacking an IP address:
   a. If Endpoint Protection uses static IP addressing pools, verify whether the static IP pool is available and also verify that each IP address assigned to the IP pool is not already in use by another VM or host. Remove the IP addresses in question from the pool, or manually free up the IP addresses in question.
   b. If Endpoint Protection uses DHCP or dynamic IP addressing, determine if the DHCP server is correctly configured.

   vSphere 6.x supports VIB downloads over port 443 (instead of port 80). This port is opened and closed dynamically. The intermediate devices between the ESXi hosts and VMware vCenter must allow traffic using this port.
3. If the NSX Manager reports that installing Endpoint Protection fails for each cluster see Installing NSX Guest Introspection services (MUX VIB) on the ESXi host fails in VMware NSX for vSphere 6.x (2135278)

Solution

Endpoint Protection should have a valid IP address allocated to it and show as Up and Running in the Service Deployment window.

Endpoint Protection Service Fails with Error

NSX Manager is unable to deploy Endpoint Protection and an alarm error is displayed.

Problem

Attempting to install Endpoint Protection (EPP) will cause the alarm error “Installation of deployment unit failed, please check if ovf/vib URLs are accessible. This error occurs quickly after attempting installation.

Cause

The NSX Manager is unable to deploy or install a EPP VM and no VM is deployed.

Solution

1. Verify if ESX Agent Manager (EAM) is up and running. After logging in to the NSX Manager, navigate to **Home** → **Alarms** tab, the installation status of the EAM connection appears as Not Ready after the NSX Manager or EAM service is restarted. Click **Resolve**.
2. If EAM is not running, or if the EAM page has issues loading and rendering, restart EAM service.
3. Check EAM logs for authentication issues. See Network port requirements for VMware NSX for vSphere 6.x (2079386) and ESX Agent Manager solution user fails to log in after replacing the vCenter Server certificates in vSphere 6.0 (2112577).
4. Verify that all required ports are open between ESXi hosts, vCenter Server and NSX Manager, and are not being blocked by a Firewall. See TCP and UDP Ports required to access VMware vCenter Server, VMware ESXi and ESX hosts, and other network components (1012382).

   vSphere 6.x supports VIB downloads over port 443 (instead of port 80). This port is opened and closed dynamically. The intermediate devices between the ESXi hosts and VMware vCenter must allow traffic using this port.

East-West Network Security - Chaining Third-party Services

After partners register network services such as Intrusion Detection System or Intrusion Protection System (IDS/IPS) with NSX, as an administrator you can configure network services to introspect east-west traffic moving between VMs on an on-premises data center.

Prerequisites

- Partners must register services with NSX.
- Prepare clusters of ESXi hosts must be prepared as NSX transport nodes by applying transport node profiles.

Note

- Service VMs are only supported on ESXi hosts.
- NSX only protects guest VMs running on ESXi hosts.

Key Concepts of East-West Network Protection

Traffic flowing between Guest VMs on an on-premises data center is protected by third-party services provided by partners. There are a few concepts that aid your understanding of the workflow.

- Service: Partners register services with NSX . A service represents the security functionality offered by the partner, service deployment details such as OVF URL of service VMs, point to attach the service, state of the service. When a notification is generated for a service, NSX notifies the partner after a time interval of 30 seconds.

- Vendor Template: It consists of functionality that a service can perform on a network traffic. Partners define vendor templates. For example, a vendor template can provide a network operation service such as tunneling with IPSec service.
- Service Profile: Is an instance of a vendor template. An NSX administrator can create a service profile to be consumed by service VMs.
- Guest VM: a source or destination of traffic in the network. The incoming or outgoing traffic is introspected by a service chain defined for a rule running east-west network services.
- Service VM: A VM that runs the OVA or OVF appliance specified by a service. It is connected over the service plane to receive redirected traffic.
- Service Instance: Is created when a service is deployed on a host. Each service instance has a corresponding service VM.
- Service Segment: A segment of a service plane that is associated to a transport zone. Each service attachment is separated from other service attachments and from the regular L2 or L3 network segments provided by NSX. The service plane manages service attachments.
- Service Manager: Is the partner service manager that points to a set of services.
- Service Chain: Is a logical sequence of service profiles defined by an administrator. Service profiles introspect network traffic in the order defined in the service chain. For example, the first service profile is firewall, second service profile is monitor, and so on. Service chains can specify different sequence of service profiles for different directions of traffic (egress/ ingress).
- Redirection Policy: Ensures that traffic classified for a specific service chain is redirected to that service chain. It is based on traffic patterns that match NSX security group and a service chain. All traffic matching the pattern is redirected along the service chain.
- Service Path: Is a sequence of service VMs that implement the service profiles of a service chain. An administrator defines the service chain, which consists of a pre-defined order of service profiles. NSX generates multiple service paths from a service chain based on the number, and locations of guest VMs and service VMs. It selects the optimum service path for the traffic flow to be introspected. Each service path is identified by a Service Path Index (SPI) and each hop along a path has a unique Service Index (SI).

NSX Requirements for East-West Traffic

For East-West Network Introspection, create a service segment and an overlay transport zone. However, you can back all other segments or logical switches on a VLAN transport zone.

East-West Network Introspection is applied to an entire NSX deployment. You can deploy the service at a cluster-level or on a per-host basis.

Multiple deployment methods are supported. One of them is host-based deployment. The type of deployment decides where service VMs run for a particular service. However, irrespective of the type of deployment, service VMs can be accessed by all East-West Network Introspection workloads. For example, a workload running on cluster A can use a service VM running on cluster B if there is no better alternative. So, picking a cluster-based deployment does not limit East-West Network Introspection to that cluster.

Even if you plan a deployment using only VLAN-backed segments, East-West traffic passes through overlay transport zones and overlay-backed segments. East-West Network Introspection is applied to all segments in the topology, whether they are backed by overlay or VLAN transport zones.

Requirements for East-West Network Introspection

- Ensure the transport nodes that host guest VMs and service VMs are configured with an overlay transport zone. An overlay transport zone is a requirement to use East-West Network Introspection on all the transport nodes in the system.
- Create an overlay-backed service segment that will be used by East-West Network Introspection service.
- All the segments must be backed by the same host switch on each host.
- If a guest VM running on an ESXi host is connected to a VLAN segment but that ESXi host is not configured to an overlay transport zone, then traffic destined to a service VM is disrupted. Such a configuration can also cause traffic to be routed to a black hole.

vMotion of Guest VMs

During a vMotion, the guest VM can be successfully migrated to another host only if the destination host is configured with an overlay transport zone and there is a single host switch. However, if there is no overlay transport zone where the service segment is created or if there are multiple host switches configured, then the virtual NIC of the guest VM goes into disconnected state even after vMotion.

Unsupported environments

- A few transport nodes are configured for VLAN transport zone, while the remaining hosts are configured for VLAN and GENEVE (overlay) transport zones. Ensure all transport nodes are configured for both VLAN and GENEVE (overlay) transport zones.
- Traffic exiting out of a guest VM virtual NIC carries .1q VLAN tag.
- Trunk port (which can carry multiple VLANs from guest VM) backed guest VMs.
- Any topology involving multiple host switches does not support east-west network introspection.

An overlay-backed (GENEVE-backed) segment is provisioned for internal use by East-West Network Introspection. On the NSX Manager UI, go to **Security → Network Introspection Settings → Service Segment**.

High-Level Tasks for East-West Network Security

Follow these steps to set up network security for east-west traffic.

Table 16-12. List of Tasks to Configure East-West Network Introspection

Workflow Tasks	Persona	Implementation
Register Service	Partner	Only API
Register Vendor Template	Partner	Only API
Register Service Manager	Partner	Only API
Deploy a Service for East-West Traffic Introspection	Administrator	API and NSX Manager UI
Add a Service Profile	Administrator	API and NSX Manager UI
Add a Service Chain	Administrator	API and NSX Manager UI
Add Redirection Rules for East-West Traffic	Administrator	API and NSX Manager UI

Deploy a Service for East-West Traffic Introspection

After partners register services, as an administrator, you must deploy an instance of the service on member hosts of a cluster.

Deploy partner service VMs that run the partner security engine on all the NSX hosts in a cluster. After you deploy the SVMs, you can create policy rules used by SVM to protect guest VMs.

Prerequisites

- All hosts are managed by a VMware vCenter.
- Partner services must be registered with NSX and are ready for deployment.
- NSX administrators can access partner services and vendor templates.
- Both the service VM and the partner service manager (console) must be able to communicate with each other at the management network level.
- Ensure only one overlay transport zone is connected to hosts that are running the partner service.
- Ensure only one service segment is used to connect partner SVM for network introspection.
- Starting with NSX 3.1, on clusters that span physical servers placed in different racks, you can override the transport node profile applied on a per-host basis.
- Starting with NSX 3.0, you must prepare clusters (cluster-based or host-based deployment methods) by applying a transport node profile.

- With NSX 2.5.x or earlier, before you deploy service VMs on each host using host-based service deployment method, configure each host of the cluster with NSX by applying a transport node profile.
- When upgrading the third-party service, the existing service will continue to be functional even if transport node profile is not applied to the cluster.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **System > Service Deployments > Deployment > Deploy Service**.

3 From the Partner Service field, select the partner service.

4 Enter the service deployment name.

5 In the Compute Manager field, select the VMware vCenter to deploy the service.

6 In the Cluster field, select the cluster where the services need to be deployed.

7 In the Data Store drop-down menu, select a data store as the repository for the service virtual machine.

8 In the Network column, click **Set** and enter the Management Network interface by choosing DHCP or static IP address type, and data network.

9 In the Service Segments field, select a service segment from the list or click the Action icon to add or edit a service segment. For every service, you can only add one service segment per overlay zone.

   A service segment is used as a service plane overlay segment where packets from guest VMs that match the classifier are sent to this service plane overlay segment and forwarded to the third-party VM for traffic introspection.

   To create a service segment:

   a Click the **+** icon next to the Service Segment field.

   b In the Service Segment dialog box, click **Add Service Segment**.

   c Enter a name, select a Transport Zone Overlay from the drop-down menu, and if applicable, select a gateway under Applied to Gateway.

   d Click **Save**.

10 In the Deployment Type field, select from one of the following deployment options. Depending upon the services registered by the partner, multiple services can be deployed as part of a single service VM.

   - Clustered: Deploys service VM on a host or hosts on one of the following types of cluster:
     - A dedicated cluster to run only service VMs.
     - A cluster that is running workload VMs. After deployment, service VMs co-exist with workload VMs.

    - An NSX Edge cluster.
  - Host Based: Deploys the service on all the hosts within a cluster.

11 In the Deployment Template field, select the template that provides attributes to protect the workload you want to run on guest VMs groups.

12 (Cluster-based deployment only) In the Clustered Deployment Count, enter the number of service VMs to deploy on the cluster.

   The VMware vCenter decides on which host to deploy the service VMs.

13 Click **Save**.

Results

After service deployment, the partner Service Manager is notified about the update.

What to do next

Know deployment details and heath status about service instances deployed on hosts. See Add a Service Profile.

Add Redirection Rules for East-West Traffic

Add rules to redirect an east-west traffic for network introspection.

Rules are defined in a policy. Policy as a concept is similar to the concept of sections in firewalls. When you add a policy, select the service chain to redirect the traffic for introspection by service profiles of the service chain.

A rule definition consists of source and destination of the traffic, introspection service, the NSX object to apply the rule to, and traffic redirection policy. After you publish the rule, NSX Manager triggers the rule when a matching traffic pattern is found. The rule begins to introspect the traffic. For example, when NSX Manager classifies a traffic flow that must be introspected, it forwards the traffic to the regular distributed firewall and then to the specified service chain in the policy. The service profiles defined in the service chain introspect the traffic for network services the partner offers. If a service profile finishes introspection without detecting any security issues in the traffic, the traffic is forwarded to the next service profile in the service chain. At the end of the service chain, the traffic is forwarded to the destination target.

All notifications are sent to the partner Service Manager and NSX.

Note By default, a rule exists even when east-west service is not configured. This default rule is not applied and is inactive. You need to create and apply the first rule after deploying east-west serviec on NSX.

Prerequisites

A service chain is available to redirect the traffic for a network introspection.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Verify the NSX Manager is in **Policy** mode.

3 Select **Security > E-W Network Introspection > Add Policy**.

A policy section is similar to a firewall section where you define rules that determine how traffics flows.

4 Select a service chain.

5 To add a policy, click **Publish**.

6 Click the  vertical ellipsis on a section and click **Add Rule**.

7 In the **Sources** column, click the edit icon, and select the source of the rule. See Add a Group for more information.

IPv4, IPv6, and multicast addresses are supported.

8 Click **Save**.

9 In the **Destinations** column, click the edit icon, and select the destination of the rule. If not defined, the destination matches any. See Add a Group for more information.

IPv4, IPv6, and multicast addresses are supported.

10 By default, the **Applied to** column is set to DFW, and the rule is applied to all workloads. You can also apply the rule or policy to selected groups.**Applied to** defines the scope of enforcement per rule, and is used mainly for optimization or resources on ESXi hosts. It helps in defining a targeted policy for specific zones and tenants, without interfering with other policy defined for other tenants and zones.

Groups consisting of only IP addresses, MAC Addresses, or Active Directory groups cannot be used in the **Applied-to** text box.

11 In the Action text box, select **Redirect** to redirect traffic along the service chain or **Do Not Redirect** not to apply network introspection on the traffic.

12 Click **Publish**.

13 To revert a published rule, select a rule and click **Revert**.

14 To add a policy, click **+ Add Policy**.

15 To clone a policy or a rule, select the policy or rule and click **Clone**.

16 To enable a rule, enable the Enable/Disable icon or select the rule and from the menu click **Enable > Enable Rule**.

17 After enabling or disabling a rule, click **Publish** to enforce the rule.

Results

Traffic going to the source is redirected to the service chain for network introspection. After service profiles in the chain introspect the traffic, it is delivered to the destination.

During deployment, it is possible that the VM group membership for a particular policy changes. NSX notifies the partner Service Manager about these updates.

Exclude Members from a Security Service

You can exclude policy groups consisting of members from being applied east-west security services.

Update the exclusion list, a list that references member groups to be excluded from the east-west service introspection policy. The excluded members are not applied with any service introspection policy.

Note

- An exclusion list does not support policy groups with IP Set, IP Addresses, or MAC Addresses as members. You can update the exclusion list from the NSX Manager
- EdgeVMs, if any, are added by system to Policy SI Exclude List through Edge_NSGroup. Removal of Edge_NSGroup might lead to traffic disruption.

An exclusion list does not support policy groups with IP Set, IP Addresses, or MAC Addresses as members. You can update the exclusion list from the NSX Manager UI or by making the following API call:

```
PUT https://<policy-mgr>/policy/api/v1/infra/settings/service-insertion/security/exclude-list

{ "members": ["/infra/domains/default/groups/grp1"], "_revision": 1 }
```

Procedure

1 Navigate to **Security > E-W Network Introspection > Actions > Exclusion List**.

A window appears listing available groups.

2 To add a user-defined group to the firewall exclusion list, click the check box next to any group. Then click **Save**.

a To create a group, click **Add Group**. See Add a Group.

b To edit a group, click the three dot menu next to a group and select **Edit**.

c To delete a group, click the three dot menu and select **Delete**.

d To display group details, click **Expand All**.

3 Click **Close**.

Get List of Service Paths

Get a list of service paths associated with a service chain, that is applied to either East-West network introspection service or Service Chaining at NSX Edge.

Know the service paths associated to a service chain ID.

Prerequisites

Procedure

- Run the API command:

```
GET https://<policy-mgr>/policy/api/v1/infra/service-chains/<service-chain-id>/service-paths

Sample Response:

    {
        "results": [{
            "service_path_id": 38,
            "service_chain_uuid": "85677e98-a3a1-4989-9c60-eeb6a04515bf",
            "service_chain_id": 2,
            "forward_path": {
                "unidir_service_path_id": 75,
                "host_cross_count": 0,
                "is_active": true,
                "in_maintenance_mode": false,
                "hops": [{
                    "mac_address": "00:50:56:83:e9:0f",
                    "vif": "4d7bfaa9-4770-435a-99de-8b9756ac025a",
                    "is_active_from_mp": true,
                    "is_active_from_ccp": true,
                    "is_active_from_dp": true,
                    "in_maintenance_mode": false,
                    "nsh_liveness_support": false,
                    "can_decrement_si": false,
                    "action": "REDIRECT"
                }]
            },
            "reverse_path": {
                "unidir_service_path_id": 76,
                "host_cross_count": 0,
                "is_active": true,
                "in_maintenance_mode": false,
                "hops": [{
                    "mac_address": "00:50:56:83:e9:0f",
                    "vif": "4d7bfaa9-4770-435a-99de-8b9756ac025a",
                    "is_active_from_mp": true,
                    "is_active_from_ccp": true,
                    "is_active_from_dp": true,
                    "in_maintenance_mode": false,
                    "nsh_liveness_support": false,
                    "can_decrement_si": false,
```

```
                    "action": "REDIRECT"
                }]
            }
        }],
        "result_count": 1
    }
```


Uninstall an East-West Traffic Introspection Service

Uninstall an east-west traffic introspection service.

As part of uninstalling an east-west service, you need to delete the east-west policy, partner service deployed, service chain, service profile, and service segment.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Verify the NSX Manager is in **Policy** mode.
3. To delete a policy, select **Security** → **East West Security** → **Network Introspection (E-W)**.
4. Select the east-west policy, click the vertical ellipses, and click **Delete Policy**.
5. Click **Publish**.
6. To delete a partner service, select **System** → **Service Deployments**.
7. Select the partner service, click the vertical ellipses and click **Delete**.
8. Click **Delete** to complete the process.
9. To delete an east-west service chain, select **Security** → **Settings** → **Network Introspection Settings** → **Service Chain**.
10. Select the service chain, click the vertical ellipses and click **Delete**.
11. To delete an east-west service profile, select **Security** → **Settings** → **Network Introspection Settings** → **Service Profile**.
12. Select the service profile, click the vertical ellipses and click **Delete**.
13. To delete an east-west service segment, select **Security** → **Settings** → **Network Introspection Settings** → **Service Segment**.
14. Select the service segment, click the vertical ellipses and click **Delete**.

15 If there are issues related to east-west service even after it was uninstalled from the NSX Manager UI, call the following API.

a Set `east_west_enabled` to **`false`** by calling the following API.

`PUT https://<nsx-manager-ip>/policy/api/v1/infra/settings/service-insertion/security/status`

```
{
    "north_south_enabled": true,
    "east_west_enabled": false,
    "resource_type": "PolicySIStatusConfiguration",
    "id": "status",
    "display_name": "status",
    "path": "/infra/settings/service-insertion/security/status",
    "relative_path": "status",
    "parent_path": "/infra",
    "unique_id": "caf620e9-405f-4533-81ab-2bd5df733364",
    "marked_for_delete": false,
    "overridden": false,
    "_create_user": "system",
    "_create_time": 1646684124017,
    "_last_modified_user": "system",
    "_last_modified_time": 1646687791212,
    "_system_owned": false,
    "_protection": "NOT_PROTECTED",
    "_revision": 0
}
```

b If tranport nodes, where the east-west service is deployed, are not connected to an overlay network, then VDS switch ports block traffic from being redirected to the east-west service. To unblock VDS switch ports, remove the extra service insertion settings on the VDS switch by running the following CLI command.

`net-dvs -u com.vmware.port.extraConfig.serviceInsertion.gvm -p <VDS_Switch_ID> nsxvswitch`

Upgrade East-West Service VM

Use API calls to upgrade an East-West service VM.

Prerequisites

Procedure

1 Use the following API to retrieve the payload of the East-West service.

`GET https://<nsx-manager>/api/v1/serviceinsertion/services/<ew_service_id>`

```
Response:
{
"functionalities": [
"NET_MON",
```

```
"NG_FW"
],q
"implementations": [
"EAST_WEST"
],
"attachment_point": [
"SERVICE_PLANE"
],
"transports": [
"NSH"
],
"on_failure_policy": "BLOCK",
"service_deployment_spec": {
"deployment_template": [
{
"name": "EW_DepTemp",
"attributes": [
{
"key": "password",
"display_name": "password",
"attribute_type": "PASSWORD",
"read_only": false
},
{
"key": "LicenseKey",
"display_name": "License",
"attribute_type": "STRING",
"read_only": false
}
]
}
],
"deployment_specs": [
{
"name": "EW_DepSpec",
"ovf_url": "http://<server-FQDN/IP>/ovfs/nsxt/EastWest/v2/EW_SI_SVM_v2.ovf",
"min_host_version": "6.5",
"host_type": "ESXI",
"service_form_factor": "MEDIUM",
"svm_version": "1.0"
}
],
"nic_metadata_list": [
{
"interface_label": "eth",
"interface_index": 0,
"interface_type": "MANAGEMENT",
"user_configurable": true
},
{
"interface_label": "eth",
"interface_index": 1,
"interface_type": "DATA1"
}
],
```

```
"svm_version": "1.0"
},
"vendor_id": "ABC_Service",
"service_capability": {
"nsh_liveness_support_enabled": true,
"can_decrement_si": false
},
"resource_type": "ServiceDefinition",
"id": "2cb0efaa-beb1-461b-be2b-d94afee98692",
"display_name": "ABC_Service",
"description": "This is East West Service Insertion",
"_create_user": "admin",
"_create_time": 1614099564224,
"_last_modified_user": "admin",
"_last_modified_time": 1614104390834,
"_system_owned": false,
"_protection": "NOT_PROTECTED",
"_revision": 1
}
```

2 Update East-West service deployment specification. The following API call adds a couple of additional deployment specifications to the payload.

`PUT https://<nsx-manager>/api/v1/serviceinsertion/services/<ew_service_id>`

```
{
"functionalities": [
"NET_MON",
"NG_FW"
],
"implementations": [
"EAST_WEST"
],
"attachment_point": [
"SERVICE_PLANE"
],
"transports": [
"NSH"
],
"on_failure_policy": "BLOCK",
"service_deployment_spec": {
"deployment_template": [
{
"name": "EW_DepTemp",
"attributes": [
{
"key": "password",
"display_name": "password",
"attribute_type": "PASSWORD",
"read_only": false
},
{
"key": "LicenseKey",
"display_name": "License",
"attribute_type": "STRING",
```

```
"read_only": false
}
]
}
],
"deployment_specs": [
{
"name": "EW_DepSpec_up2",
"ovf_url": "http://<server-FQDN/IP>/OVAs/SVM-tiny-64_3/SVM-tiny-64/SVM-tiny-64.ovf",
"min_host_version": "6.5",
"host_type": "ESXI",
"service_form_factor": "SMALL",
"svm_version": "1.0"
},
{
"name": "EW_DepSpec",
"ovf_url": "http://<server-FQDN/IP>/ovfs/nsxt/EastWest/v2/EW_SI_SVM_v2.ovf",
"min_host_version": "6.5",
"host_type": "ESXI",
"service_form_factor": "MEDIUM",
"svm_version": "1.0"
},
{
"name": "EW_DepSpec_up",
"ovf_url": "http://<server-FQDN/IP>/EW_OVF/EW_SI_SVM.ovf",
"min_host_version": "6.5",
"host_type": "ESXI",
"service_form_factor": "LARGE",
"svm_version": "1.0"
}
],
"nic_metadata_list": [
{
"interface_label": "eth",
"interface_index": 0,
"interface_type": "MANAGEMENT",
"user_configurable": true
},
{
"interface_label": "eth",
"interface_index": 1,
"interface_type": "DATA1"
}
],
"svm_version": "1.0"
},
"vendor_id": "ABC_Service",
"service_capability": {
"nsh_liveness_support_enabled": true,
"can_decrement_si": false
},
"resource_type": "ServiceDefinition",
"id": "2cb0efaa-beb1-461b-be2b-d94afee98692",
"display_name": "ABC_Service",
"description": "This is East West Service Insertion",
```

```
"_create_user": "admin",
"_create_time": 1614099564224,
"_last_modified_user": "admin",
"_last_modified_time": 1614104390834,
"_system_owned": false,
"_protection": "NOT_PROTECTED",
"_revision": 1
}
```

3 Call the following API to upgrade Service VMs belonging to East-West service deployment. NSX upgrades East-West Service VM associated with the specified service deployment.

```
POST https://<nsx-manager>/api/v1/serviceinsertion/services/
<ew_service_id>/service-deployments/<service_deployment_id>?action=upgrade
```

```
{
"deployment_spec_name": "EW_DepSpec_up2"
}
```

Results

NSX upgrades the East-West Service VM.

North-South Network Security - Inserting Third-party Service

NSX provides the functionality to insert third-party services at tier-0 or tier-1 router in the data center to redirect traffic to the third-party service for introspection. Only ESXi hosts are supported to deploy north-south service VMs.

High-Level Tasks for North-South Network Security

Follow these steps to set up network security for north-south traffic.

Table 16-13. List of Tasks to Configure North-South Network Introspection

Workflow Tasks	Persona	Implementation
Register Service with NSX	Partner	Only API
Deploy a Service for North-South Traffic Introspection	Administrator	API and NSX UI
Add Redirection Rules for North-South Traffic	Administrator	API and NSX UI

Deploy a Service for North-South Traffic Introspection

After you register a service, you must deploy an instance of the service on an NSX transport node for the service to start processing network traffic.

Deploy partner service VM at tier-0 or tier-1 logical router that acts as a gateway between the physical world and the logical network on VMware vCenter. After you deploy the SVM as a standalone service instance or an active-standby service instance, you can create redirection rules to redirect traffic to the SVM for network introspection.

Prerequisites

- All hosts are managed by a VMware vCenter.
- Partner services are registered with NSX and are ready for deployment.
- NSX administrators can access partner services.
- High Availability mode for logical router must be in active-standby mode.
- Turn on the Distributed Resource Scheduler utility.
- Ensure only one overlay transport zone is connected to hosts that are running the partner service.
- Ensure only one service segment is used to connect guest VMs for network introspection.
- Starting with NSX 3.1, on clusters that span physical servers placed in different racks, you can override the transport node profile applied on a per-host basis.
- Starting with NSX 3.0, you must prepare clusters (cluster-based or host-based deployment methods) by applying a transport node profile.
- With NSX 2.5.x or earlier, before you deploy service VMs on each host using host-based service deployment method, configure each host of the cluster with NSX by applying a transport node profile.
- When upgrading the third-party service, the existing service will continue to be functional even if transport node profile is not applied to the cluster.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **System > Service Deployments > Deployment**.

3 From the Partner Service drop-down menu, select the service that must be deployed.

4 Click **Deploy Service**. Enter details to deploy the service.

Table 16-14. Partner Service Details

Field	Description
Service Deployment Name	Enter a name to identify the service instance.
Deployment Specification	Select the form factor to deploy.
Attachment Points	Select the tier-0 or tier-1 logical router where the service instance must be deployed.
Failure Policy	Select **Allow** or **Block**.

Table 16-14. Partner Service Details (continued)

Field	Description
Network	For a deployment of the type **Active Standby**, set value to the following fields: ■ Primary Interface Network: The interface to be used by the deployed service. ■ Primary Interface IP: Enter the IP address to be used by the service instance. ■ Primary Gateway Address: Enter the gateway address. ■ Primary Subnet Mask: Enter the subnet mask. ■ Secondary Interface Network: The standby interface that is used if the primary interface is unavailable. ■ Secondary Interface IP: Enter the IP address for the standby IP that is used if the primary IP is unavailable. ■ Secondart Gateway Address: Enter the standby gateway address that is used if the primary gateway is unavailable. ■ Secondary Subnet Mask: Enter the standby subnet mask that is used if the primary subnet mask is unavailable. For a deployment of the type **Standalone**, set values to the primary interfaces.
Compute Manager	Select the registered VMware vCenter.
Datastore	Select the repository to store service instance data.
Deployment Mode	Select **Standalone** to deploy a single service instance at the tier-0 or tier-1 logical router. Select **Active Standby** to deploy a couple of service instances in active-standby mode at the tier-0 or tier-1 logical router.
Deployment Template	Select the template to be used during deployment of the service instance.

5 Click **Save**.

Results

The Service Instances tab displays the deployment progress. It might take a few minutes for deployment to finish. Verify the deployment state to ensure that the service instance is successfully deployed at the tier-0 or tier-1 logical router.

Alternatively, go to the VMware vCenter and verify the deployment status.

What to do next

Add redirection rules for north-south traffic. See Add Redirection Rules for North-South Traffic.

Add Redirection Rules for North-South Traffic

Set up redirection rules to send traffic to third-party services inserted at a Tier-0 or Tier-1 router.

Prerequisites

- Register and deploy third-party services on NSX.
- Configure Tier-0 or Tier-1 router.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Select **Security > North South Security > Network Introspection (N-S) > Add Policy**.

   A policy section is similar to a firewall section where you define rules that determine how traffics flows.
3. Set **Redirection To** field for a service instance or a service chain to a Tier-0 or Tier-1 logical router to perform network introspection of traffic flowing between source and destination entities.
4. To add a policy, click **Publish**.
5. Click the  vertical ellipsis on a section and click **Add Rule**.
6. Edit the **Source** field to add a group by defining membership criteria, static members, IP/MAC addresses, or active directory groups. Membership criteria can be defined from one of these types: Virtual Machine, Logical Switch, Logical Port, IP Set. You can select static members from one of these categories: Group, Segment, Segment Port, Virtual Network Interface, or Virtual Machine.
7. Click **Save**.
8. To add a destination group, edit the **Destination** field.
9. In the **Applied To** field, you can do one of the following:
   - For a service inserted at Tier-0 logical router, select the uplink of Tier-0 router.
   - For a service inserted at Tier-1 logical router, you do not need to select any uplinks.
10. Each rule can be enabled individually. After you enable a rule, it is applied to the traffic that matches the rule.
11. Click **Advanced Settings** to configure the traffic direction and to enable logging.
12. In the Action field, select **Redirect** to redirect traffic along the service instance or **Do Not Redirect** not to apply network introspection on the traffic.
13. Click **Publish**.
14. To revert a published rule, select a rule and click **Revert**.
15. To add a policy, click **+ Add Policy**.

16 To clone a policy or a rule, select the policy or rule and click **Clone**.

17 To enable a rule, enable the Enable/Disable icon or select the rule and from the menu click **Enable > Enable Rule**.

18 After enabling or deactivating a rule, to enforce the rule, click **Publish**.

Results

Based on the actions set, NSX redirects north-south traffic to the service instance for network introspection.

Traffic on NSX Edge nodes is redirected to a service path for traffic introspection. After the service path introspects traffic, packets are sent to their original destination.

Starting in NSX 3.2, north-south traffic redirected to a service chain can use multiple service paths for load balancing. NSX selects one of the optimal service paths currently available to serve each new traffic flow. Each flow is pinned to a single path. Different flows can use different paths based on round robin policy. North-South service chaining can use a maximum number of 16 service paths.

Uninstall a North-South Traffic Introspection Service

Uninstall a north-south traffic introspection service.

Delete a policy and a partner service that was deployed for north-south introspection service.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Verify the NSX Manager is in **Policy** mode.

3 To delete a policy, select **Security → North South Security → Network Introspection (N-S)**.

4 Select the policy, click the vertical ellipses, and click **Delete Policy**.

5 Click **Publish**.

6 To delete a partner service, **System → Service Deployments**.

7 Select the service, click the vertical ellipses and click **Delete**.

Update Service Insertion Status

Update status of third-party services, such as East-West network introspection or North-South network introspection.

You can enable or disable North-South or East-West service insertion configuration status on the system. By default, service insertion status is disabled until a service is created.

- Update the status of service. Run the following API command.

```
PUT https://<policy-mgr>/policy/api/v1/infra/settings/service-insertion/security/status
    {
      "north_south_enabled": false,
      "resource_type": "PolicySIStatusConfiguration",
      "_revision": 1
    }
Example Response:
{
    "resource_type": "PolicySIStatusConfiguration",
    "id": "si-status",
    "display_name": "/infra/settings/service-insertion/security/status",
    "path": "/infra/settings/service-insertion/security/status",
    "relative_path": "si-status",
    "north_south_enabled": false,
    "east_west_enabled": true,
    "_create_user": "system",
    "_create_time": 1554274751846,
    "_last_modified_user": "admin",
    "_last_modified_time": 1554275071006,
    "_system_owned": false,
    "_protection": "NOT_PROTECTED",
    "_revision": 2
}
```

Upgrade North-South Service VM

Use API calls to upgrade an North-South Service VM.

Prerequisites

Procedure

1 Use the following API to retrieve the payload of the North-South service deployment.

`GET https://<nsx-manager>/api/v1/serviceinsertion/services/<ns_service_id>`

```
{
"functionalities": [
"NG_FW"
],
"implementations": [
"NORTH_SOUTH"
],
"attachment_point": [
"TIER1_LR",
"TIER0_LR"
],
"transports": [
"L2_BRIDGE"
```

```
],
"on_failure_policy": "BLOCK",
"service_deployment_spec": {
"deployment_template": [
{
"name": "NS_DepTemp",
"attributes": [
{
"key": "LicenseKey",
"display_name": "License",
"attribute_type": "STRING",
"read_only": false
}
]
}
],
"deployment_specs": [
{
"name": "NS_DepSpec",
"ovf_url": "http://<server-FQDN/IP>/ovfs/nsxt/NorthSouth/v1/PA-VM-NST-8.1.3-c15.ovf",
"min_host_version": "6.5",
"host_type": "ESXI",
"service_form_factor": "MEDIUM",
"svm_version": "1.0"
}
],
"nic_metadata_list": [
{
"interface_label": "eth",
"interface_index": 2,
"interface_type": "DATA2"
},
{
"interface_label": "eth",
"interface_index": 0,
"interface_type": "MANAGEMENT",
"user_configurable": true
},
{
"interface_label": "eth",
"interface_index": 1,
"interface_type": "DATA1"
},
{
"interface_label": "eth",
"interface_index": 4,
"interface_type": "HA2"
},
{
"interface_label": "eth",
"interface_index": 3,
"interface_type": "HA1"
}
]
},
```

```
"vendor_id": "ABC_Service",
"resource_type": "ServiceDefinition",
"id": "9ea92ea7-a6b0-4640-9c28-5b364713e8c0",
"display_name": "NS Service_automation",
"description": "This service is inserted at T0 & T1 router and it provides advanced
security",
"_create_user": "admin",
"_create_time": 1614099580424,
"_last_modified_user": "admin",
"_last_modified_time": 1614099580424,
"_system_owned": false,
"_protection": "NOT_PROTECTED",
"_revision": 0
}
```

2 Update the service deployment specification. In this call, deployment sepcs are updated with **NS_Deployment_Specfication_11** and **NS_Deployment_Specfication_12**.

```
PUT https://<nsx-manager>/api/v1/serviceinsertion/services/<ns_service_id>
```

```
{
"functionalities": [
"NG_FW"
],
"implementations": [
"NORTH_SOUTH"
],
"attachment_point": [
"TIER1_LR",
"TIER0_LR"
],
"transports": [
"L2_BRIDGE"
],
"on_failure_policy": "BLOCK",
"service_deployment_spec": {
"deployment_template": [
{
"name": "NS_DepTemp",
"attributes": [
{
"key": "LicenseKey",
"display_name": "License",
"attribute_type": "STRING",
"read_only": false
}
]
}
],
"deployment_specs": [
{
"name": "NS_Deployment_Specfication_11",
"ovf_url": "http://<server-FQDN/IP>/ovfs/nsxt/NorthSouth/v1/PA-VM-NST-8.1.3-c15.ovf",
"host_type": "ESXI"
},
```

```
{
"name": "NS_Deployment_Specfication_12",
"ovf_url": "http://<server-FQDN/IP>/ovfs/nsxt/NorthSouth/v1/PA-VM-NST-8.1.3-c15.ovf",
"host_type": "ESXI"
}
],
"nic_metadata_list": [
{
"interface_label": "eth",
"interface_index": 2,
"interface_type": "DATA2"
},
{
"interface_label": "eth",
"interface_index": 0,
"interface_type": "MANAGEMENT",
"user_configurable": true
},
{
"interface_label": "eth",
"interface_index": 1,
"interface_type": "DATA1"
},
{
"interface_label": "eth",
"interface_index": 4,
"interface_type": "HA2"
},
{
"interface_label": "eth",
"interface_index": 3,
"interface_type": "HA1"
}
]
},
"vendor_id": "ABC_Service",
"resource_type": "ServiceDefinition",
"id": "9ea92ea7-a6b0-4640-9c28-5b364713e8c0",
"display_name": "NS Service_automation",
"description": "This service is inserted at T0 & T1 router and it provides advanced
security",
"_create_user": "admin",
"_create_time": 1614099580424,
"_last_modified_user": "admin",
"_last_modified_time": 1614099580424,
"_system_owned": false,
"_protection": "NOT_PROTECTED",
"_revision": 0
}
```

3 Update the Tier-0 service instance deployment specification name to **NS_Deployment_Specfication_12** . To update the specification, use the response returned by the GET call (earlier step) and feed the same in the following PATCH call with `deployment_spec_name`: **NS_Deployment_Specfication_12**. NSX upgrades the Tier-0 Service VM.

```
PATCH https://<nsx-manager>/policy/api/v1/infra/tier-0s/
<tier0_LR_id>/locale-services/<tier0_localeservice_id>/service-instances/
<tier0_serviceinstance_id>
```


```
{
"context_id": "b75193d1-137c-45c2-83b3-8f9d7795ddca",
"compute_id": "b75193d1-137c-45c2-83b3-8f9d7795ddca:domain-c16",
"storage_id": "datastore-13",
"primary_interface_network": "/infra/tier-1s/TIER-1-LR-1/segments/SI_SEGMENT",
"primary_interface_mgmt_ip": "10.161.154.2",
"secondary_interface_network": "/infra/tier-1s/TIER-1-LR-1/segments/SI_SEGMENT",
"secondary_interface_mgmt_ip": "10.161.154.3",
"deployment_spec_name": "NS_Deployment_Specfication_12",
"failure_policy": "BLOCK",
"deployment_template_name": "NS_DepTemp",
"attributes": [
{
"key": "LicenseKey",
"display_name": "License",
"value": "adklfjakldfjlkajfkl",
"attribute_type": "STRING",
"read_only": false
}
],
"resource_type": "PolicyServiceInstance",
"id": "PAN-1",
"display_name": "PAN-1",
"path": "/infra/tier-0s/Tier0-LR-1/locale-services/tier0localservices/service-instances/
PAN-1",
"relative_path": "PAN-1",
"parent_path": "/infra/tier-0s/Tier0-LR-1/locale-services/tier0localservices",
"unique_id": "5500a8e4-255b-4131-9da2-9a9a2be21b32",
"marked_for_delete": false,
"overridden": false,
"partner_service_name": "NS Service_automation",
"deployment_mode": "ACTIVE_STANDBY",
"transport_type": "L2_BRIDGE",
"_create_user": "admin",
"_create_time": 1614099590522,
"_last_modified_user": "admin",
"_last_modified_time": 1614099590527,
"_system_owned": false,
"_protection": "NOT_PROTECTED",
"_revision": 0
}
```

4 Update the service instance deployment specification of the Tier-1 Logical Router 1 to **`NS_Deployment_Specfication_12`** . For this you can use the response returned by GET call and feed the same to PATCH call with `deployment_spec_name`: **`NS_Deployment_Specfication_12`**. NSX upgrades the Tier-1 gateway Service VM.

`PATCH https://<nsx-manager>/policy/api/v1/infra/tier-1s/<tier1_LR_id>/locale-services/<tier1_LR1_localeservice_id>/service-instances/<tier1_LR1_serviceinstance_id>`


```
{
"context_id": "b75193d1-137c-45c2-83b3-8f9d7795ddca",
"compute_id": "b75193d1-137c-45c2-83b3-8f9d7795ddca:domain-c16",
"storage_id": "datastore-13",
"primary_interface_network": "/infra/tier-1s/TIER-1-LR-1/segments/SI_SEGMENT",
"primary_interface_mgmt_ip": "10.161.154.4",
"secondary_interface_network": "/infra/tier-1s/TIER-1-LR-1/segments/SI_SEGMENT",
"secondary_interface_mgmt_ip": "10.161.154.5",
"deployment_spec_name": "NS_Deployment_Specfication_12",
"failure_policy": "BLOCK",
"deployment_template_name": "NS_DepTemp",
"attributes": [
{
"key": "LicenseKey",
"display_name": "License",
"value": "adklfjakldfjlkajfkl",
"attribute_type": "STRING",
"read_only": false
}
],
"resource_type": "PolicyServiceInstance",
"id": "PAN-1-tier-1",
"display_name": "PAN-1-tier-1",
"path": "/infra/tier-1s/TIER-1-LR-1/locale-services/1-policyconnectivity-1284/service-instances/PAN-1-tier-1",
"relative_path": "PAN-1-tier-1",
"parent_path": "/infra/tier-1s/TIER-1-LR-1/locale-services/1-policyconnectivity-1284",
"unique_id": "4bd48611-3257-4895-8295-c49fdfa644b3",
"marked_for_delete": false,
"overridden": false,
"partner_service_name": "NS Service_automation",
"deployment_mode": "ACTIVE_STANDBY",
"transport_type": "L2_BRIDGE",
"_create_user": "admin",
"_create_time": 1614099595487,
"_last_modified_user": "admin",
"_last_modified_time": 1614099595489,
"_system_owned": false,
"_protection": "NOT_PROTECTED",
"_revision": 0
}
```

5 Update the service instance deployment specification of Tier-1 Logical Route 2 to **NS_Deployment_Specfication_12**. To update the deployment specification, use the response returned by GET call (earlier step) and paste the same in the following PATCH call with `deployment_spec_name`: **NS_Deployment_Specfication_12**. NSX upgrades the Tier-1 Logical Router 2 Service VM.

```
PATCH https://<nsx-manager>/policy/api/v1/
infra/tier-1s/<tier1_LR2_id>/locale-services/<tier1_LR2_localeservice_id>/
service-instances/<tier1_LR2_serviceinstance_id>
```


```
{
"context_id": "b75193d1-137c-45c2-83b3-8f9d7795ddca",
"compute_id": "b75193d1-137c-45c2-83b3-8f9d7795ddca:domain-c16",
"storage_id": "datastore-13",
"primary_interface_network": "/infra/tier-1s/TIER-1-LR-1/segments/SI_SEGMENT",
"primary_interface_mgmt_ip": "10.161.154.6",
"secondary_interface_network": "/infra/tier-1s/TIER-1-LR-1/segments/SI_SEGMENT",
"secondary_interface_mgmt_ip": "10.161.154.7",
"deployment_spec_name": "NS_Deployment_Specfication_12",
"failure_policy": "BLOCK",
"deployment_template_name": "NS_DepTemp",
"attributes": [
{
"key": "LicenseKey",
"display_name": "License",
"value": "adklfjakldfjlkajfkl",
"attribute_type": "STRING",
"read_only": false
}
],
"resource_type": "PolicyServiceInstance",
"id": "PAN-2-tier-1",
"display_name": "PAN-2-tier-1",
"path": "/infra/tier-1s/TIER-1-LR-2/locale-services/2-policyconnectivity-1359/service-
instances/PAN-2-tier-1",
"relative_path": "PAN-2-tier-1",
"parent_path": "/infra/tier-1s/TIER-1-LR-2/locale-services/2-policyconnectivity-1359",
"unique_id": "eb9e61e7-edc3-4c29-93cc-d3c4eeb4a6c7",
"marked_for_delete": false,
"overridden": false,
"partner_service_name": "NS Service_automation",
"deployment_mode": "ACTIVE_STANDBY",
"transport_type": "L2_BRIDGE",
"_create_user": "admin",
"_create_time": 1614099599703,
"_last_modified_user": "admin",
"_last_modified_time": 1614099599705,
"_system_owned": false,
"_protection": "NOT_PROTECTED",
"_revision": 0
}
```

Results

NSX upgrades the North-South Service VM.

Network Introspection Settings

This section contains settings to configure network introspection.

Service insertion (Network Introspection) support only occurs when an NSX Federation environment has a Global Manager (GM) deployed under the following conditions:

- All service-insertion related configuration such as partner service registration, deployment and consumption, is done from a Local Manager (LM).
- Only objects configured on the LM are used with service insertion. This includes groups, segments, and any other constructs. Service insertion cannot be applied to workloads connected to a stretched/global segment defined from the GM, or any segment connected to a logical router created from the GM. Groups created from the Global Manager should not be used within service insertion redirection polices.

Add a Service Segment

Add a service segment when you configure east-west network introspection on the overlay network or north-south network introspection, where you want to redirect packets from the uplink of an NSX Edge to the service chain.

Prerequisites

- If you are configuring north-south service chaining to redirect packets from the uplink of an NSX Edge to the service chain, create a Tier-0 and or Tier-1 gateway. The segment is later connected to the Tier-0 and or Tier-1 gateway.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Click **Security > Settings > Network Introspection Settings > Service Segment > Add Service Segment**.
3. Click **Add Service Segment**.
4. In the **Name** field, enter a name for the segment.
5. In the **Transport Zone (Overlay)** field, select an overlay transport zone that is associated to the segment.
6. In the **Connected To** field, do one of the following:
   - Leave the field blank if you are configuring east-west network introspection to protect guest VMs by third-party security vendors.
   - Select a Tier-0 or Tier-1 gateway if you are configuring a north-south service chaining to redirect packets from the uplink of an NSX Edge to the service chain.

7 Click **Save**.

Results

The **Status** column displays the status of the service segment.

Add a Service Profile

A service profile is an instance of a partner vendor template. Administrators can customize attributes of a vendor template to create an instance of the template.

Note You can create multiple service profile for a single vendor. For example, the service profile set for the forward path provides IDS protection, whereas the service profile set for the reverse path supports IPS protection. However, a single service profile can be set for both forward and reverse path.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Navigate to **Security > Network Introspection > Service Profiles > Add Service Profile**.

3 From the Partner Service drop-down field, select a service. You can create a service profile for the selected service.

4 Enter the service profile name and select the vendor template.

5 The Redirection Action field inherits functionality from the vendor template. For example, if COPY is the functionality provided by the vendor template, then by default the redirection action when you create a service profile is COPY.

6 (Optional) Define any tags to filter out and manage service profiles.

7 Click **Save**.

Results

A new service profile is created for the partner service.

What to do next

Add a service chain. See Add a Service Chain.

Add a Service Chain

A service chain is a logical sequence of service profiles defined by the network administrator.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Verify the NSX Manager is in **Policy** mode.

3 Navigate to

Security > Network Introspection > Service Chains > Add Chain

4 Enter the service chain name.

5 In the Service Segments field, select the service segment to which you want to apply the service chain.

A service segment is a segment of service plane that connects multiple service VMs of an overlay transport zone. Each service VM in the service chain is separate from another service VM and L2 and L3 network segments run by NSX. The service plane controls access to service VMs.

6 To set the forward path, click the **Set Forward Path field** and click **Add Profile in Sequence**.

7 Add the first profile in the service chain and click **Add**.

8 To specify the next service profile, click **Add Profile in Sequence** and enter details.

You can also rearrange the profile order by using the Up and Down arrow icons.

9 Click **Save** to finish adding a forward path for the service chain.

10 In the Reverse Path column, select **Inverse Forward Path** for the service plane to use the service profile you set for the forward path.

11 To set a new service profile for the reverse path, click **Set Reverse Path** and add a service profile.

12 Click **Save** to finish adding a reverse path for the service chain.

13 In the Failure Policy field,

- Select **Allow** to send traffic to the destination VM when the service VM fails. Service VM failure is detected by the liveness detection mechanism which can be enabled only by partners.
- Select **Block** to not send traffic to the destination VM when the service VM fails.

14 Click **Save**.

Results

After adding a service chain, the partner Service Manager is notified about the update.

What to do next

Create a redirection rule to introspect east-west network traffic. See Add Redirection Rules for East-West Traffic.

NSX IDS/IPS and NSX Malware Prevention

NSX Intrusion Detection and Prevention Service (IDS/IPS) monitors east-west traffic and north-south traffic to detect malicious traffic patterns by comparing the traffic against a known set of

intrusion detection signatures. NSX Malware Prevention extracts files from the east-west traffic and north-south traffic and analyzes these files for malicious behavior.

These two security features are a part of the VMware NSX® Advanced Threat Prevention solution.

Getting Started with NSX IDS/IPS and NSX Malware Prevention

Read the topics in this section to obtain an overview of NSX IDS/IPS and NSX Malware Prevention features. Understand the system requirements, terminologies used, and complete the prerequisites tasks to prepare your data center for using these two features.

Overview of NSX IDS/IPS and NSX Malware Prevention

The objective of NSX Intrusion Detection and Prevention Service (IDS/IPS) is to monitor network traffic on the hosts and edges for malicious activity by comparing the traffic against a known set of signatures. The objective of NSX Malware Prevention is to extract files from the network traffic on the hosts and edges and analyze these files for malicious behavior.

Overview of NSX Intrusion Detection and Prevention Service

NSX IDS/IPS monitors network traffic on a host for suspicious activity by comparing the traffic against signatures. A signature specifies a pattern for a type of network intrusion that needs to be detected and reported. Whenever a matching traffic pattern to a signature is found, a predefined action is taken, such as generating an alert or blocking the traffic from reaching its destination.

Implementation of IDS is carried through the following methods:

- Knowledge-based signatures: Knowledge-based signatures incorporate specific knowledge or pattern that corresponds to a known type of attack. In this approach, IDS attempts to detect intrusions based on already known malicious instruction sequences specified in signatures. Thus, knowledge-based signatures are limited to attacks that are already known and cannot cover targeted or zero-day threats.
- Behavior-based detection: Behaviour-based detection attempts to identify anomalous behavior by pinpointing interesting events that are different or unusual compared to a baseline or normal traffic.

  These events are called informational or info and consists of events that pinpoint unusual activities in a network that are not necessarily malicious but can provide valuable information when investigating a breach. Signatures are bundled together with custom detection logic that can be updated without having to recompile or modify the IDS engine. Behavior-based detection introduces a new IDS intrusion severity level as 'suspicious'.

NSX supports IDS/IPS capability on both Distributed and Gateway firewall. NSX IDS/IPS on Gateway Firewall feature only on tier-1 gateways.

Overview of NSX Malware Prevention

NSX Malware Prevention can detect and prevent known malicious files and unknown malicious files. Unknown malicious files are also referred to as zero-day threats. To detect malware, NSX Malware Prevention uses a combination of the following techniques :

- Hash-based detection of known malicious files
- Local analysis of unknown files
- Cloud analysis of unknown files

The supported maximum file size limit for malware analysis is 64 MB.

NSX Distributed Malware Prevention

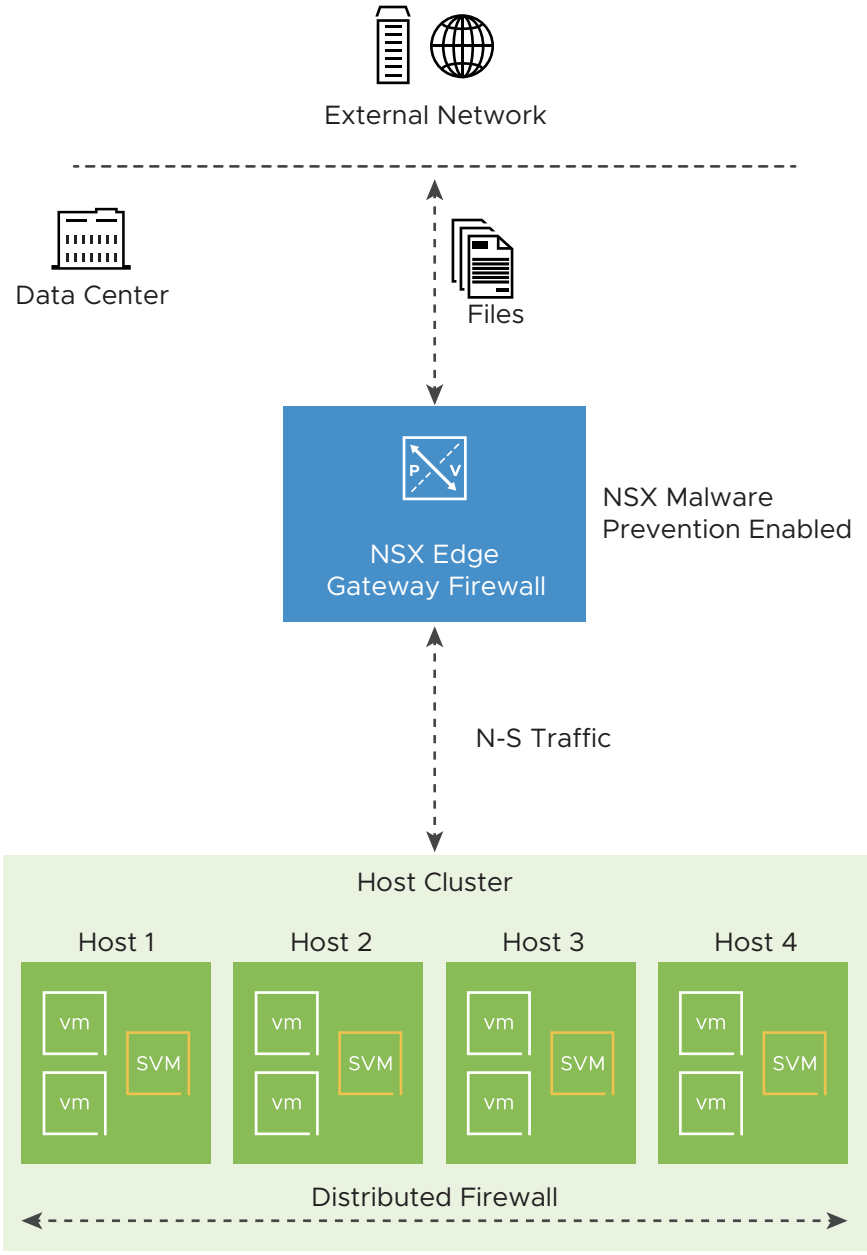
Malware detection and prevention on the Distributed Firewall is supported for both Windows and Linux guest endpoints (VMs), which are running on vSphere host clusters that are prepared for NSX.

Local analysis and cloud analysis of all categories of malware files are supported. To view the list of supported file categories, see File Categories Supported for NSX Malware Prevention.

NSX Malware Prevention on Gateway Firewall

Only detection of malware is supported on the Gateway Firewall. Local analysis and cloud analysis of all categories of malware files are supported. To view the list of supported file categories, see the hyperlinked topic that is mentioned in the NSX Distributed Malware Prevention section.

Figure 16-8. Concept diagram of NSX Malware Prevention



NSX Malware Prevention on E-W Traffic

On the north-south traffic, the NSX Malware Prevention feature uses the IDS/IPS engine on the NSX Edges to extract or intercept the files that are entering the data center. On the east-west traffic, this feature uses the capabilities of the NSX Guest Introspection (GI) platform. If the file bypasses scrutiny on the NSX Edge and reaches the host, the file is extracted by the GI thin agent on guest VMs.

To detect and prevent malware on guest VMs, you must install the NSX Guest Introspection thin agent on guest VMs, and deploy the NSX Distributed Malware Prevention service on vSphere host clusters that are prepared for NSX. When this service is deployed, a service virtual machine (SVM) is installed on each host of the vSphere cluster and NSX Malware Prevention is enabled on the host cluster.

NSX Guest Introspection thin agent drivers for Windows are included with VMware Tools. To view the VMware Tools versions that are supported for your version of NSX, see the VMware Product Interoperability Matrix. To view the list of supported Windows guest operating systems for a specific VMware Tools version, see the release notes for that version in the VMware Tools documentation.

Note Windows 11 and Windows 2022 guest OS versions are supported with VMware Tools 12.0.6 or later.

The Guest Introspection thin agent for Linux is available as part of the operating system specific packages (OSPs). The packages are hosted on VMware packages portal. Installing open-vm-tools or VM Tools is not required for Linux. To view the list of supported Linux guest operating system versions, see the Prerequisites section in Install the Guest Introspection Thin Agent for Anti-virus on Linux Virtual Machines.

NSX Malware Prevention File Events

File events are generated when files are extracted by the IDS engine on the NSX Edges in the north-south traffic and by the NSX Guest Introspection agent on the virtual machine endpoints in the distributed east-west traffic.

NSX Malware Prevention feature inspects the extracted files to determine whether they are benign, malicious, or suspicious. Each unique inspection of a file is counted as a single file event in NSX. In other words, a file event refers to a unique file inspection.

For information about monitoring the NSX Malware Prevention file events by using the UI, see Monitoring File Events.

For information about monitoring the file events by using the NSX Malware Prevention file event APIs, see the documentation on the VMware Developer Documentation portal.

System Requirements for NSX IDS/IPS and NSX Malware Prevention

NSX environment must meet specific license and software requirements to use NSX IDS/IPS and NSX Malware Prevention features.

Requirements for NSX Intrusion Detection and Prevention Service

License Requirements

For NSX Intrusion Detection and Prevention Service, the Threat Prevention license is required. To read more about NSX security licenses, see the *Security Licenses* section in License Types.

Prerequisites

For Gateway IDS/IPS, NSX Edge VMs with at least the Large form factor must be deployed.

Requirements for NSX Malware Prevention

License Requirements

For NSX Malware Prevention feature, the Advanced Threat Prevention license is required.

For example:

- NSX Distributed Firewall with Advanced Threat Prevention license
- NSX Gateway Firewall with Advanced Threat Prevention license

To read more about NSX security licenses, see the *Security Licenses* section in License Types.

Prerequisites

The following prerequisites are common to both Distributed NSX Malware Prevention and Gateway NSX Malware Prevention:

- NSX Application Platform must be deployed and NSX Malware Prevention feature must be activated on the platform.
- Internet access is required even when files are not sent to the cloud for a detailed analysis. For more information, see the Notes section after this bulleted list.
- NSX Manager nodes and vSphere hosts must have connectivity to the NSX Application Platform for NSX Malware Prevention to function properly.
- Minimum supported vSphere version is 6.7
- Minimum supported VMware Tools version is 11.2.5

Notes (IP Access to External Sites)

NSX Malware Prevention feature requires Internet access to download the latest signatures and to send files for cloud analysis. The following communication is done on HTTPS:

- From NSX Application Platform (K8s worker IP address) or HTTP proxy if the platform is configured with proxy.
- To NSX Advanced Threat Prevention cloud service:
  - nsx.lastline.com
  - nsx.west.us.lastline.com if you selected "Malware Cloud Region = United States" during installation

- nsx.nl.emea.lastline.com if you selected “Malware Cloud Region = European Union” during installation

The following prerequisites apply only to Distributed NSX Malware Prevention:

- Windows VMs must have VMware Tools with NSX File Introspection driver.
- Linux VMs must have the File Introspection driver for Linux.
- On each vSphere host, service virtual machine (SVM) deployment requires following reources:
  - 4 vCPU
  - 6 GB RAM
  - 80 GB Disk space
- Web server is required to deploy the SVM.
- vSphere host clusters must be configured with a transport node profile.

The following prerequisite applies only to Gateway NSX Malware Prevention:

- NSX Edge VMs must be deployed with Extra Large form factor, or use bare metal edge nodes.

Note NSX Malware Prevention feature is not supported on Public Cloud Gateways.

Preparing the Data Center for NSX IDS/IPS and NSX Malware Prevention

You can set up NSX IDS/IPS and NSX Malware Prevention features in your NSX environment only when your data center uses an appropriate license.

For information about licenses that are required to run the NSX Advanced Threat Prevention solution, see the *Security Licenses* section in License Types.

Preparing the data center for NSX Intrusion Detection/Prevention and NSX Malware Prevention involves multiple steps. To do these steps, you can use the **IDS/IPS & Malware Prevention Setup** wizard.

The setup wizard is like an onboarding process that guides you through a sequence of steps to prepare the data center for these two security features. To run this wizard, navigate to **Security > IDS/IPS & Malware Prevention**.

If NSX detects that appropriate licenses are not added, the page displays the following text:

```
IDS/IPS & Malware Prevention is not supported with current license.
```

If NSX detects that appropriate licenses are added, the page displays the **Start Setup** and **Skip Setup** buttons.

To begin the setup wizard, click **Start Setup**. Follow the on-screen instructions and this documentation to complete the steps in the wizard.

- If you want to save your progress at any stage and exit the wizard, click **Back to Main Page**. Later, you can continue the setup from where you left off.
- If you want to reset the setup wizard, and start again from the beginning, click **Cancel**. Canceling the setup removes the selections you made in the wizard, but it does not remove any deployments that you completed in the wizard. For example, if you completed the deployment of the NSX Application Platform and the NSX Malware Prevention service virtual machine on host clusters before resetting the wizard, these deployments are retained.
- If you do not want to use the setup wizard and prefer setting up the two security features on your own later, click **Skip Setup**. NSX Manager does not show this wizard again. Later, you can navigate to **Security > IDS/IPS & Malware Prevention > Settings** and set up the data center for both the features. For information about using the **IDS/IPS & Malware Prevention Settings** page, see Configuring NSX IDS/IPS and NSX Malware Prevention Settings.

By default, all the check boxes in the IDS/IPS and Malware Prevention feature cards are selected for setup. You can edit the selections, if required. When you are ready to proceed, click **Next**. Your selections determine the tabs that are shown in the wizard, as explained in the following table.

Note NSX Application Platform is a prerequisite for NSX Malware Prevention, but not for NSX IDS/IPS.

Selected Features	Tabs Shown
IDS/IPS on east-west traffic or IDS/IPS on north-south traffic	Configure NSX Proxy Manage Signatures Enable Nodes
Malware Prevention only on east-west traffic	Configure NSX Proxy Deploy NSX Application Platform Deploy Service VM
Malware Prevention only on north-south traffic	Configure NSX Proxy Deploy NSX Application Platform Enable Nodes
Malware Prevention on both east-west traffic and north-south traffic	Configure NSX Proxy Deploy NSX Application Platform Deploy Service VM Enable Nodes
All features selected	All five tabs in the wizard are shown

Configure NSX Proxy Server for Internet Connectivity

NSX Malware Prevention can work only when your NSX is connected to the Internet. NSX IDS/IPS can work in a network without Internet connectivity, but you will have to manually update the IDS/IPS signatures.

Click the **Go to NSX Proxy Server** link and specify the following settings:

- Scheme (HTTP or HTTPS)
- IP address of the host
- Port number
- User name and password

Deploy NSX Application Platform

NSX Malware Prevention requires certain microservices to be deployed in the NSX Application Platform. You must first deploy the NSX Application Platform, and then activate the NSX Malware Prevention feature. After this feature is activated, the microservices that are required for NSX Malware Prevention get deployed in the platform.

To summarize, you must perform the following tasks in the given order:

1 Deploy NSX Application Platform

2 Activate NSX Malware Prevention

Note Versioning of the NSX Malware Prevention feature in the NSX Application Platform matches the NSX Application Platform version number, and not the NSX product version number.

Deploy Service Virtual Machine

For east-west traffic in the data center, you must deploy the NSX Distributed Malware Prevention service on vSphere host clusters that are prepared for NSX. When this service is deployed, a service virtual machine (SVM) is installed on each host of the vSphere cluster and NSX Malware Prevention is enabled on the host cluster.

A donut chart on this page shows the number of host clusters in the data center where the NSX Distributed Malware Prevention service is deployed and not deployed.

For detailed instructions about deploying the NSX Distributed Malware Prevention service on a host cluster, see Deploy the NSX Distributed Malware Prevention Service.

After the service deployment is done on the host clusters, return to this page in the wizard, and click **Next** to continue.

Note High availability is not supported for the service virtual machine of NSX Distributed Malware Prevention service.

Manage Signatures

When Internet connectivity is configured in your data center, NSX Manager checks for availability of new intrusion detection signatures on the cloud every 20 minutes, by default. When a new update is available, a banner is displayed on the page with an **Update Now** link.

If the data center does not have an Internet connectivity, you can manually download the IDS signature bundle (`.zip`) file, and then upload the file to NSX Manager. For detailed instructions, see Offline Downloading and Uploading NSX Intrusion Detection Signatures.

Signature Management

Signature management tasks are optional. If needed, you can do them later by navigating to **Security > IDS/IPS & Malware Prevention > Settings > IDS/IPS**.

- To view signature version or to add another version of the signatures in addition to the default, click **View and Change**.

  Currently, two versions of signatures are maintained. Whenever there is a change in the version commit identification number, a new version is downloaded.

- To automatically download intrusion detection signatures from the cloud and apply them to the hosts and edges in the data center, turn on the **Auto Update** toggle.

  When this option is turned off, the automatic download of signatures stops. You can manually download the IDS signature bundle (`.zip`) file, and then upload the file to NSX Manager.

- To view status of signature download on transport nodes, click the link in **Status** field.

- To globally exclude specific signatures or to change their action to alert, drop, or reject, click **View and Manage Signature Set**.

  Select an **Action** for the signature, and click **Save**. The changes done in global signature management settings are applicable to all IDS/IPS profiles. However, if you update the signature settings in an IDS/IPS profile, the profile settings take precedence.

  The following table explains the meaning of each signature action.

Action	Description
Alert	An alert is generated and no automatic preventive action is taken.
Drop	An alert is generated and the offending packets are dropped.
Reject	An alert is generated and the offending packets are dropped. For TCP flows, a TCP reset packet is generated by IDS and sent to the source and destination of the connection. For other protocols, an ICMP-error packet is sent to the source and destination of the connection.

Enable Nodes for IDS/IPS and Malware Prevention

In the **Activate Hosts & Clusters for East-West Traffic** section, do the following configurations:

- Turn on NSX IDS/IPS on the standalone ESXi hosts.
- Select the ESXi host clusters where you want to turn on NSX IDS/IPS on the east-west traffic.

- If the NSX Distributed Malware Prevention service is not already deployed on ESXi host clusters, click the **Defined in Service VM deployment** link in the **Malware Prevention** column. For instructions about deploying the NSX Distributed Malware Prevention service on a host cluster, see Deploy the NSX Distributed Malware Prevention Service.

Note

- Do not enable NSX Distributed IDS/IPS in an environment that is using Distributed Load Balancer. NSX does not support IDS/IPS with a Distributed Load Balancer.
- For NSX Distributed IDS/IPS to work, Distributed Firewall (DFW) must be enabled. If traffic is blocked by a DFW rule, then IDS/IPS cannot see the traffic.

In the **Activate Gateways for North-South Traffic** section, do the following configurations:

- Select the tier-1 gateways where you want to turn on NSX IDS/IPS on the north-south traffic.
- Select the tier-1 gateways where you want to turn on NSX Malware Prevention on the north-south traffic.

Important On the north-south traffic, NSX supports:

- NSX Malware Prevention feature only on tier-1 gateways.
- NSX IDS/IPS on Gateway Firewall feature only on tier-1 gateways.

Configuring NSX IDS/IPS and NSX Malware Prevention Settings

If you have skipped the **IDS/IPS and Malware Prevention Setup** wizard without configuring any settings, or if you have skipped the wizard midway during the configuration process, you can continue the configuration process from the **IDS/IPS & Malware Prevention Settings** page.

To open this page in the NSX Manager UI, navigate to **Security > IDS/IPS & Malware Prevention > Settings**.

The configuration settings are grouped into three tab pages:

- Shared
- IDS/IPS
- Malware Prevention

Shared Settings

As the name suggests, these settings are common to NSX IDS/IPS and NSX Malware Prevention.

Configure Internet Proxy Server

NSX Malware Prevention can work only when your NSX is connected to the Internet. NSX IDS/IPS can work in a network without Internet connectivity, but you will have to manually update the IDS/IPS signatures.

Click the **Internet Proxy Server** link and specify the following settings:

- Scheme (HTTP or HTTPS)
- IP address of the host
- Port number
- User name and password

Define Scope for Malware Prevention and IDS/IPS Deployment

In the **Activate Hosts & Clusters for East-West Traffic** section, do the following configurations:

- Turn on NSX IDS/IPS on the standalone ESXi hosts.
- Select the ESXi host clusters where you want to turn on NSX IDS/IPS on the east-west traffic.
- If the NSX Distributed Malware Prevention service is not already deployed on ESXi host clusters, click the **Defined in Service VM deployment** link in the **Malware Prevention** column. For instructions about deploying the NSX Distributed Malware Prevention service on a host cluster, see Deploy the NSX Distributed Malware Prevention Service.

In the **Activate Gateways for North-South Traffic** section, do the following configurations:

- Select the tier-1 gateways where you want to turn on NSX IDS/IPS on the north-south traffic.
- Select the tier-1 gateways where you want to turn on NSX Malware Prevention on the north-south traffic.

Important On the north-south traffic, NSX supports:

- NSX Malware Prevention feature only on tier-1 gateways.
- NSX IDS/IPS on Gateway Firewall feature only on tier-1 gateways.

IDS/IPS Settings

When Internet connectivity is configured in your data center, NSX Manager checks for availability of new intrusion detection signatures on the cloud every 20 minutes, by default. When a new update is available, a banner is displayed on the page with an **Update Now** link.

If the data center does not have an Internet connectivity, you can manually download the IDS signature bundle (`.zip`) file, and then upload the file to NSX Manager. For detailed instructions, see Offline Downloading and Uploading NSX Intrusion Detection Signatures.

You can perform the following signature management tasks on this page:

- To view signature version or to add another version of the signatures in addition to the default, click **View and Change**.

Currently, two versions of signatures are maintained. Whenever there is a change in the version commit identification number, a new version is downloaded.

- To automatically download intrusion detection signatures from the cloud and apply them to the hosts and edges in the data center, turn on the **Auto Update** toggle.

  When this option is turned off, the automatic download of signatures stops. You can manually download the IDS signature bundle (`.zip`) file, and then upload the file to NSX Manager.

- To view status of signature download on transport nodes, click the link in **Status** field.
- To globally exclude specific signatures or to change their action to alert, drop, or reject, click **View and Manage Signature Set**.

  Select an **Action** for the signature, and click **Save**. The changes done in global signature management settings are applicable to all IDS/IPS profiles. However, if you update the signature settings in an IDS/IPS profile, the profile settings take precedence.

  The following table explains the meaning of each signature action.

Action	Description
Alert	An alert is generated and no automatic preventive action is taken.
Drop	An alert is generated and the offending packets are dropped.
Reject	An alert is generated and the offending packets are dropped. For TCP flows, a TCP reset packet is generated by IDS and sent to the source and destination of the connection. For other protocols, an ICMP-error packet is sent to the source and destination of the connection.

You can also manage the following advanced settings:

- To send IDS/IPS events to external syslog consumers, turn on the **Syslog** toggle.
- To configure whether excess traffic should be dropped or bypassed by the IDS/IPS engine in case of oversubscription, click the appropriate option in the **Oversubscription** field.

Malware Prevention Settings

NSX Malware Prevention requires certain microservices to be deployed in the NSX Application Platform.

If the NSX Application Platform is not deployed in your data center, this page displays the following title:

```
Malware Prevention is not deployed yet.
```

Do the following steps:

1 Read the on-screen text and click **Go to NSX Application Platform**.

2 Before proceeding with the platform deployment, read the NSX Application Platform Deployment Checklist in the *Deploying and Managing the VMware NSX Application Platform* publication at https://docs.vmware.com/en/VMware-NSX-T-Data-Center/index.html. From the left navigation pane at this link, expand the product version, and then click the publication name.

3 Deploy the NSX Application Platform. For more details, see the *Deploying and Managing the VMware NSX Application Platform* publication.

4 Activate NSX Malware Prevention feature on the platform.

After the NSX Malware Prevention feature is activated on the NSX Application Platform, the **Malware Prevention** settings page displays the **Allowlist** section. You might have to refresh the page a few times to view this section.

Allowlist

Using the NSX Manager UI or API, you can override or suppress the verdict of the file that NSX has computed. This overridden file verdict takes precedence over the NSX computed verdict. The Allowlist table table displays all the files with suppressed verdict. This table is initially empty. When you start monitoring the file events in your data center by using the **Malware Prevention** dashboard and suppress file verdicts based on your specific security requirements, the suppressed files are added to the Allowlist table.

To learn more about overriding file verdicts, see Add a File to the Allowlist.

Terminology Used in NSX Malware Prevention

Familiarize yourself with the key terminologies that are used in NSX Malware Prevention.

Cloud File Analysis

Cloud file analysis is done by the NSX Advanced Threat Prevention service that is running in the cloud. It involves a detailed analysis of unknown files by using the following techniques to detect whether the file is benign, malicious, or suspicious:

- NSX Malware Prevention sandboxing and behavioral analysis
- Statistical algorithms
- Artificial intelligence and machine learning
- Deep content inspection

NSX sends unknown files over a secure connection to the cloud only when you opt for cloud file analysis in your Malware Prevention security profile.

File Event

An event that is generated when a file is extracted or intercepted from the data path traffic on an NSX Edge or a Guest VM on the host. On an NSX Edge, the file is extracted by the NSX IDPS engine, and on a Guest VM, the file is extracted by the NSX File Introspection driver in the Guest Introspection (GI) thin agent.

Local File Analysis

Local file analysis is done inside the NSX on NSX Edge Transport Nodes and ESXi Host Transport Nodes that are activated for NSX Malware Prevention. It involves a lightweight scanning of unknown files against a known set of file hashes to detect whether the file is benign, malicious, or suspicious.

Malware Class

It is the type of threat. Examples of malware class are virus, trojan horse, worm, adware, ransomware, spyware, and so on.

Malware Family

It is a name that identifies a specific group of malware files, which typically originate from the same source code or developed by the same malware authors. Examples of malware families are valyria, darkside, and so on.

Reputation

Threat information about a file, URL, or other artifacts that provides details about the file, URL.

For example, reputation of a file can include the following details:

- Name of the file publisher
- Is the file signed (Yes or No)
- The signing authority of the file
- Reputation category of the file (malware, suspect, trusted)
- Malware class to which the file belongs to. For example, Trojan horse, backdoor, adware, and so on.

File reputation details are stored in the cloud and accessible to all the NSX customers.

Threat Score

It denotes the degree of risk or malicious intent that is associated with the file. A high threat score indicates a greater amount of risk, and the reverse.

Verdict

NSX Malware Prevention reports a decision about the files, which are intercepted in the data center either on the NSX Edges (north-south traffic) or on the guest VMs (east-west traffic). The decision about the file is called a verdict. Verdict can be one of the following values.

Value	Description
Benign	The file is good or safe to be downloaded.
Trusted	The file is trusted based on its behavior.
Highly Trusted	The file is from a highly trusted source, for example, Microsoft, Apple, Adobe, and so on.

Value	Description
Malicious	The file is harmful or a threat to the data center.
Suspicious	The file is potentially harmful or unwanted.
Unknown	The file is not known to NSX and therefore no decision is available for the file.
Uninspected	This file is not inspected by NSX Malware Prevention because you had earlier suppressed or allowlisted the file.

Zero-day Threat

A threat that is not seen in NSX before and which does not match any of the known malware signatures.

File Categories Supported for NSX Malware Prevention

NSX Malware Prevention supports multiple file categories for both local file analysis and cloud file analysis.

The following file categories are supported:

- Executable
- Document
- Script
- Archive
- Data
- Media
- Other

On the Distributed Firewall, NSX Malware Prevention feature is supported for both Windows and Linux guest endpoints (VMs).

On the Distributed Firewall, NSX Malware Prevention supports both detection and prevention of malware. However, on the Gateway Firewall, only detection of malware is supported.

The sections that follow later in this topic contain examples of supported file extensions for each file category. These examples only serve as a reference and should not be interpreted as the complete list of supported file extensions for each file category. Other file extensions for these listed file categories are also supported for analysis. The maximum file size limit is 64 MB.

Executable Files

The following table lists examples of supported file extensions that belong to the executable file category.

File Extensions	Description
.exe	Portable Executable/MS-DOS executable Self-extracting (SFX) executable
.elf	Executable and Linkable Format (ELF) executable
.msi	Microsoft installer
.lnk	Microsoft Windows shortcut
.dll	Microsoft Windows library
.sys	Microsoft Windows driver
.cpl, .ocx, .pif	Other Microsoft file formats that might contain executable content
.class	Compiled Java class code
.com	COM executable for DOS EICAR test virus

Document Files

The following table lists examples of supported file extensions that belong to the document file category.

File Extensions	Description
.doc, .docx	Microsoft Office Word document
.xls, .xlsx	Microsoft Office Excel document
.xlt	Microsoft Office Excel template
.xlam	Microsoft Office Excel add-in with macros
.xlsm	Microsoft Office Excel document with macros
.xlsb	Microsoft Office Excel document with macros and saved in a binary format
.xltx	Microsoft Office Excel spreadsheet template
.xltm	Microsoft Office Excel spreadsheet template with macros
.ppt, .pptx	Microsoft Office Powerpoint document
.ppsx	Microsoft Office Powerpoint slideshow
.pot, .potx	Microsoft Office Powerpoint template
.docm	Microsoft Office Word document, Office Open XML format, with macros
.pptm	Microsoft Office Powerpoint document with macros
.ppsm	Microsoft Office Powerpoint slideshow with macros
.potm	Microsoft Office Powerpoint presentation template with macros
.dot, .dotx	Microsoft Office Word document template
.dotm	Microsoft Office Word document template, Office Open XML format with macros
.xps	Microsoft XML paper specification document

File Extensions	Description
.odp, .ods, .odt, .otg, .otp, .ott, .odg	Open Office or LibreOffice document formats
.oxps	Open XML paper specification format document
.pdf	PDF document
.wpd	WordPerfect document
.pub	Microsoft Publisher document
.rtf	Rich text format document
.xml	XML-based Microsoft Office Excel document, pre-Office2007 XML-based Microsoft Office Powerpoint presentation, pre-Office2007 XML-based Microsoft Office Word document, pre-Office2007
.xdp	Adobe XML data package format
.xsl	eXtensible stylesheet language for XML file

Script Files

The following table lists examples of supported file extensions that belong to the script file category.

File Extensions	Description
.hta	HTML application (HTA)
.vba	Visual Basic for applications
.vbs	Visual Basic script
.vbe	Visual Basic encoded script
.bat, .cmd	Batch script
.js	JavaScript Analysis of Javascript files is supported only in the context of file transfers and not in the context of web traffic.
.jse	Jscript encoded script
.pl, .pm	Perl script
.psm1, .psd1, .ps1	Powershell script module Powershell data file Powershell script
.py	Python script
.sh, .command	Shell script Terminal command file
.wsf	Windows script

Archive Files

The following table lists examples of supported file extensions that belong to the archive file category.

File Extensions	Description
.ace	WinAce compressed file
.tbz2, .tbz, .bz2, .bz	TAR archive files compressed with Linux-based Bzip and Bzip2 data compressors
.cab	Microsoft Windows cabinet archive file
.diagcab	Microsoft diagnostic cabinet archive file
.tgz, .gz	TAR achive file compressed with Gnu Zip
.jar	Java archive file
.war	Java Web application archive
.lzh, .lha	Archive file compressed using Lempel-Ziv and Haruyasu (LZH) compression algorithm
.lzma	Files compressed with Lempel-Ziv-Markov chain Algorithm (LZMA) compression
.nupkg	NuGet package file
.udf	Universal disk format
.iso	Disc image file format based on ISO-9660 standard
.rar	Files compressed with RAR compression
.tar	Tape archive file
.xz, .txz	XZ compressed TAR file
.zip	Zip archive file
.7z	7-zip archive file
.eml	RFC2822-formatted email message file

Data Files

The following table lists examples of supported file extensions that belong to the Data file category.

File Extensions	Description
.csv	Comma-separated values data file
.iqy	Internet query data file
.sylk, .slk	Symbolic link data file
.pcapng, .pcap	Packet capture file (tcpdump)
.settingcontent-ms	Microsoft content-settings data file

Media Files

Only Macromedia Flash data file (`.swf`) is supported.

Other Files

The following table lists examples of supported file extensions that do not belong to any of the preceding file categories.

File Extensions	Description
`.website`	Website file
`.url`	Internet shortcut file referenced by Web browsers
`.htm, .html`	HTML document Analysis of HTML files is supported only in the context of file transfers and not in the context of web traffic. The context is detected using content-disposition headers for HTTP, and is always true for other protocols, such as FTP, SMB.
`.xar, .pkg`	XAR archive data For malware detection, these files are analyzed directly without extracting them. Therefore, they are not classified as archive files.

Activate NSX Malware Prevention

After you meet the system requirements and the following prerequisites, you can use the NSX Manager UI to activate the NSX Malware Prevention feature.

Prerequisites

- NSX Application Platform must be successfully deployed on your NSX 3.2 or later environment and in a good state. See the *Deploying and Managing the VMware NSX Application Platform* documentation for details.
- You must have NSX **Enterprise Administrator** privileges.
- A valid NSX Data Center edition license, or, starting with NSX 4.0.1.1, an NSX Advanced or NSX Enterprise Plus license that supports Advanced Threat Prevention, is in effect for your NSX Manager session. For information about licenses that are required to run the NSX Advanced Threat Prevention solution, see the *Security Licenses* section in License Types.

Procedure

1. From your browser, log in with **Enterprise Administrator** privileges to an NSX Manager at `https://nsx-manager-ip-address`.
2. In the NSX Manager UI, select **System** and in the Configuration section, select **NSX Application Platform** .
3. Navigate to the Features section, locate the NSX Malware Prevention feature card, and click **Activate** or anywhere in the card.
4. In the NSX Malware Prevention activation window, select one of the available cloud regions from which you can access the NSX Advanced Threat Prevention cloud service.

   The system uses the NSX Advanced Threat Prevention cloud service to perform deeper analysis on detected threat events, perform alert correlation and visualization, and fetch

periodic updates on those detected threats events. The regions displayed are the ones available to your NSX location. If you previously activated the VMware NSX® Network Detection and Response™ feature, the cloud region selected for that feature is the same region that is used for the NSX Malware Prevention feature.

5 Click **Run Prechecks**.

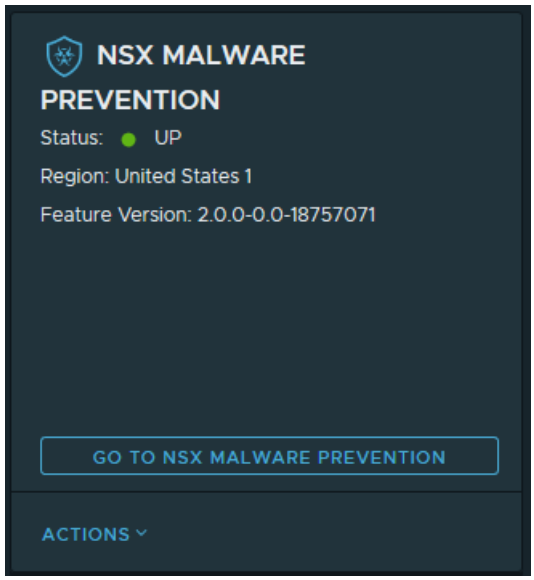
This precheck process can take some time as the system validates that the minimum license requirement is met and that it is eligible for use with the NSX Advanced Threat Prevention cloud service. The system also validates that the selected cloud region is reachable.

6 Click **Activate**.

This step can take some time.

Results

If the feature activation is successful, the NSX Malware Prevention feature card, similar to the following image, displays the Status as `UP`, information about the cloud region to which it is connected, and the version of the feature that is being used.



Offline Downloading and Uploading NSX Intrusion Detection Signatures

If Internet connectivity is not configured in your NSX, you can use APIs to manually download the NSX intrusion detection signature bundle (`.zip`) file, and then upload the signature bundle to NSX Manager. Perform the following steps to download signatures in an offline mode and upload them on NSX.

Step 1: Register NSX to the Cloud Service

Use the following API to register NSX to the cloud service. Before starting any communication with the cloud service, you must use this API to register to the cloud service. Send all licenses and you will be provided with necessary permission. If the license key is valid, the API generates and returns `client_id` and `client_secret`. The information about the license is stored in the cloud. `Client_secret` is used as the request for the Authentication API. If the client has previously registered, but does not have access to `client_id` and `client_secret`, the client has to re-register using the same API.

URI Path:

```
POST https://api.nsx-sec-prod.com/1.0/auth/register
```

Body:

```
{
"license_keys":["XXXXX-XXXXX-XXXXX-XXXXX"],
"device_type":"NSX-Idps-Offline-Download",
"client_id": "client_username"
}
```

Response:

```
{"client_id":"client_username",
"client_secret": "Y54+V/
rCpEm50x5HAUIzH6aXtTq7s97wCA2QqZ8VyrtFQjrJih7h0alItdQn02T46EJVnSMZWTseragTFScrtIwsiPSX7APQIC7M
xAYZ0BoAWvW2akMxyZKyzbYZjeROb/C2QchehC8GFiFNpwqiAcQjrQHwHGdttX4zTQ="
}
```

Step 2: Authenticate NSX to the Cloud Service

Use the following API to authenticate NSX to the cloud service. This API call authenticates the client using the `client_id` and `client_secret`. The API generates `access_token` to use in the headers of requests to IDS Signatures APIs. The token is valid for 60 minutes. If the token expires, the client has to reauthenticate using the `client_id` and `client_secret`.

URI Path:

```
POST https://api.nsx-sec-prod.com/1.0/auth/authenticate
```

Body:

```
{"client_id":"client_username",
"client_secret": "Y54+V/
rCpEm50x5HAUIzH6aXtTq7s97wCA2QqZ8VyrtFQjrJih7h0alItdQn02T46EJVnSMZWTseragTFScrtIwsiPSX7APQIC7M
xAYZ0BoAWvW2akMxyZKyzbYZjeROb/C2QchehC8GFiFNpwqiAcQjrQHwHGdttX4zTQ="
}
```

Response:

```
{
    "access_token":
"eyJhbGciOiJIUzUxMiJ9.eyJqdGkiOiI3ZjMwN2VhMmQwN2IyZjJjYzM5ZmU5NjJjNmZhNDFhMGZlMTk4YjMyMzU4OGU5
NGU5NzE3NmNmNzk0YWU1YjdjLTJkYWY2MmE3LTYxMzctNGJiNS05NzJlLTE0NjZhMGNkYmU3MCIsInN1YiI6IjdmMzA3ZW
EyZDA3YjJmMmNjMzlmZTk2MmM2ZmE0MWEwZmUxOThiMzIzNTg4ZTk0ZTk3MTc2Y2Y3OTRhZTViN2MtMmRhZjYyYTctNjEz
Ny00YmI1LTk3MmUtMTQ2NmEwY2RiZTcwIiwiZXhwIjoxNTU1NTUyMjk0LCJpYXQiOjE1NTU1NDg2OTR9.x4U75GShDLMhy
iyUO2B9HIi1Adonzx3Smo01qRhvXuErQSpE_Kxq3rzg1_IIyvoy3SJwwDhSh8KECtGW50eCPg",
    "token_type": "bearer",
    "expires_in": 3600,
    "scope": "[idps_scope]"
}
```

Step 3: Retrieve Link of the Signature Bundle (Zip) File

Use the following API to retrieve link of the signature bundle file. NSX Cloud downloads the latest signatures every 24 hours, and saves the signatures in a ZIP file. This API returns a ZIP file link in the response. Copy and paste the link in your browser and download the ZIP file.

URI Path:

```
GET https://api.prod.nsxti.vmware.com/2.0/intrusion-services/signatures
```

In the Headers tab, the Authorization key will have the `access_token` value from the authenticate API response.

```
Authorization:
eyJhbGciOiJIUzUxMiJ9.eyJqdGkiOiI3ZjMwN2VhMmQwN2IyZjJjYzM5ZmU5NjJjNmZhNDFhMGZlMTk4YjMyMzU4OGU5N
GU5NzE3NmNmNzk0YWU1YjdjLTJkYWY2MmE3LTYxMzctNGJiNS05NzJlLTE0NjZhMGNkYmU3MCIsInN1YiI6IjdmMzA3ZWE
yZDA3YjJmMmNjMzlmZTk2MmM2ZmE0MWEwZmUxOThiMzIzNTg4ZTk0ZTk3MTc2Y2Y3OTRhZTViN2MtMmRhZjYyYTctNjEzN
y00YmI1LTk3MmUtMTQ2NmEwY2RiZTcwIiwiZXhwIjoxNTU1NTUyMjk0LCJpYXQiOjE1NTU1NDg2OTR9.x4U75GShDLMhyi
yUO2B9HIi1Adonzx3Smo01qRhvXuErQSpE_Kxq3rzg1_IIyvoy3SJwwDhSh8KECtGW50eCPg
```

Response:

```
{
    "signatures_url": "https://ncs-idps-us-west-2-prod-signatures.s3.us-west-2.amazonaws.com/
a07fe284ff70dc67194f2e7cf1a8
                        178d69570528.zip?X-Amz-Security-
Token=IQoJb3JpZ2luX2VjENf%2F%2F%2F%2F%2F%2F%2F%2F%2FwEaCXVzLXd

lc3QtMSJHMEUCIG1UYbzfBxOsm1lvdj1k36LPyoPota0L4CSOBMXgKGhmAiEA%2BQC1K4Gr7VCRiBM4ZTH2WbP2rvIp0qf
HfG

lOx0ChGc4q6wEIHxABGgw1MTAwMTM3MTE1NTMiDA4H4ir7eJl779wWWirIAdLIx1uAukLwnhmlgLmydZhW7ZExe%2BamDk
RU7K

T46ZS93mC1CQeL00D2rjBYbCBiG1mzNILPuQ2EyxmqxhEOzFYimXDDBER4pmv8%2BbKnDWPg08RNTqpD%2BAMicYNP7Wlp
xeZw

YxeoBFruCDA2l3eXS6XNv3Ot6T2a%2Bk4rMKHtZyFkzZREIIcQlPg7Ej5q62EvvMFQdo8TyZxFpMJBc4IeG0h1k6QZU1Jl
krq
```

```
2RYKit5WwLD%2BQKJrEdf4A0YctLbMCDbNbprrUcCADMKyclu8FOuABuK90a%2BvnA%2FJFYiJ32eJl%2Bdt0YRbTnRyvl
MuS

UHxjNAdyrFxnkPyF80%2FQLYLVDRWUDatyAo10s3C0pzYN%2FvMKsumExy6FIcv%2FOLoO8Y9RaMOTnUfeugpr6YsqMCH0
pUR

4dIVDYOi1hldNCf1XD74xMJSdnviaxY4vXD4bBDKPnRFFhOxLTRFAWVlMNDYggLh3pV3rXdPnIwgFTrF7CmZGJAQBBKqax
zP

MVZ2TQBABmjxoRqCBip8Y662Tbjth7iM2V522LMVonM6Tysf16ls6QU9IC6WqjdOdei5yazK%2Fr9g%3D&X-Amz-
Algorithm
                        =AWS4-HMAC-SHA256&X-Amz-Date=20191202T222034Z&X-Amz-
SignedHeaders=host&X-Amz-Expires=3599&X-Amz-
                        Credential=ASIAXNPZPUTA6A7V7P4X%2F20191202%2Fus-
west-1%2Fs3%2Faws4_request&X-Amz-

Signature=d85ca4aef6abe22062e2693acacf823f0a4fc51d1dc07cda8dec93d619050f5e",
    "signature_version": "3",
    "sha256_checksum":"d0b7a20eb9b57c9abd90f95c6a29af5f2056fea4a01eb9c8a30a87e0d0fd43a5",
    "last_updated": "2020-03-01T23:59:59Z",
    "version_name": "IDS/IPSSignature.3.2020-03-01T23:59:59Z"
   }
```

Step 4: Upload the Signature Bundle to NSX Manager

- Method 1: Upload using NSX Manager UI

  To upload the file from NSX Manager UI, navigate to **Security > IDS/IPS & Malware Prevention > Settings > IDS/IPS**, and click **Upload IDS/IPS Signatures**. Browse the saved signature ZIP file and upload the file.

- Method 2: Upload using an NSX API

  To upload the file using the NSX API, use the following API.

  ```
  POST https://<mgr-ip>/policy/api/v1/infra/settings/firewall/security/intrusion-services/
  signatures?action=upload_signatures
  ```

Error Code Handling for Authentication API

This is an example authentication API error response:

```
{
"error_code":100101,
"error_message":"XXXXX"
}
```

- If you received an error code from 100101-100150, re-register with the **same** client id.
- If you received an error code from 100151-100200, re-register with a **different** client id.

Adding Security Profiles

Security profiles include IDS/IPS profile and Malware Prevention profile. To enforce NSX IDS/IPS and NSX Malware Prevention security protection in your data center, you must attach security profiles to Distributed Firewall rules and Gateway Firewall rules.

Add an NSX IDS/IPS Profile

You can create NSX IDS/IPS profiles to group signatures, which can then be applied to selected applications. You can create 100 custom profiles in addition to the default profile.

The default IDS profile includes critical severities and cannot be edited.

Procedure

1. Navigate to **Security > IDS/IPS & Malware Prevention > Profiles**.
2. Click **Add Profile**.
3. Enter a name for this profile.
4. (Optional) Enter a description for the profile and add tags.
5. Select the required **Intrusion Severities** that you want to include in the profile.
6. (Optional) Filter signatures to include in the profile by **Attack Types**, **CVSS**, **Attack Targets**, and **Products Affected**.
7. To change the action on a specific signature, click **Manage signatures for this profile** and in the **Action** column, select the appropriate action.
8. (Optional) To view only user-modified signatures, click **Show only User modified signatures** toggle button.
9. Click **Save** to create the profile.

What to do next

Create IDS rules.

Add a Malware Prevention Profile

A Malware Prevention profile determines the file categories that you want to analyze for malware, and whether you want NSX to send the files to the cloud for a detailed analysis.

You can use either the default Malware Prevention profile in your firewall rules or add new profiles depending on the requirements of your security policies. In the profile, you can select the file categories that NSX Malware Prevention should capture and analyze for malicious behavior. File analysis is done locally on NSX Host Transport Nodes and NSX Edge Transport Nodes that are activated for NSX Malware Prevention. If you opt to send the files to the cloud, a detailed file analysis is also done in the cloud.

When you apply the profile to Distributed Malware Prevention rules, NSX Malware Prevention analyzes the files that are intercepted or captured on the Host Transport Nodes. When you apply the profile to Gateway Malware Prevention rules, NSX Malware Prevention analyzes the files that are intercepted or captured on the Edge Transport Nodes.

You can add multiple Malware Prevention profiles with different configurations and use separate profiles in the Distributed Malware Prevention firewall rules and Gateway Malware Prevention firewall rules. You can use a different profile in the firewall rules of each tier-1 gateway that you have activated for NSX Malware Prevention. For example, let us say you have two profiles: A and B. In profile A configuration, you choose not to send the files to the cloud for analysis, whereas in profile B, you choose to send the files to the cloud for analysis. You use profile A for Distributed Malware Prevention rules and profile B for Gateway Malware Prevention rules.

Caution You must observe caution while using different or inconsistent Malware Prevention profile configurations for Distributed Malware Prevention rules and Gateway Malware Prevention rules, or while using different profile configurations on each tier-1 gateway that is activated for NSX Malware Prevention. Using different profile configurations might result in NSX returning a different verdict for a file depending on where the file is intercepted (Host Transport Node or Edge Transport Node). Cloud file analysis performs deep content inspection, including sandboxing and machine learning techniques. This deep content inspection can detect malware behaviors with greater accuracy. Local file analysis does not have sufficient resources to perform such a deep analysis and hence might yield less precise results.

You can attach only a single Malware Prevention profile to a firewall rule at a time. However, a single Malware Prevention profile can be attached to multiple Distributed Malware Prevention rules and Gateway Malware Prevention rules simultaneously, if required.

Prerequisites

Set up your NSX for NSX Malware Prevention.

For detailed instructions, see Preparing the Data Center for NSX IDS/IPS and NSX Malware Prevention.

Procedure

1. From your browser, log in with **admin** privileges to an NSX Manager at https://*nsx-manager-ip-address*.
2. Navigate to **Security > IDS/IPS & Malware Prevention > Profiles > Malware Prevention**.
3. Click **Add Profile**.
4. Enter a name for the profile.
5. (Optional) Enter a description for the profile and add tags.
6. Select the file categories to include for local file analysis and cloud file analysis. By default, all categories are selected.

7 (Optional) Deselect the **Send files to NSX Advanced Threat Prevention cloud service** check box.

By default, cloud file analysis is selected.

8 Click **Save**.

Results

Malware Prevention profile is saved and the **Status** column shows `Successful`.

What to do next

Attach this profile to Gateway Malware Prevention rules or Distributed Malware Prevention rules, or both, depending on the requirements of your security policies.

Using NSX IDS/IPS and NSX Malware Prevention on a Distributed Firewall

You can use the NSX IDS/IPS feature to detect malicious traffic patterns in the distributed east-west traffic, and use the NSX Malware Prevention feature to detect malicious files in the distributed east-west traffic.

Workflow for NSX Distributed IDS/IPS

Perform the following steps to use NSX Distributed IDS/IPS.

1 Set up NSX Proxy Server for Internet Connectivity. NSX IDS/IPS can work in a network without Internet connectivity, but you will need to manually update the IDS/IPS signatures. For more information, see Preparing the Data Center for NSX IDS/IPS and NSX Malware Prevention.

2 Download latest signature set and configure signature settings: Download the latest signature set if you have not selected automatic download option and configure actions for signatures. For more information, see Preparing the Data Center for NSX IDS/IPS and NSX Malware Prevention.

3 Enable nodes for NSX Distributed IDS/IPS: Select hosts on which you want to enable IDS/IPS. For more information, see Preparing the Data Center for NSX IDS/IPS and NSX Malware Prevention.

Note

- Do not enable NSX Distributed IDS/IPS in an environment that is using Distributed Load Balancer. NSX does not support IDS/IPS with a Distributed Load Balancer.
- For NSX Distributed IDS/IPS to work, Distributed Firewall (DFW) must be enabled. If traffic is blocked by a DFW rule, then IDS/IPS cannot see the traffic.

4 Create IDS/IPS profiles: Create profiles to group signatures. For more information, see Add an NSX IDS/IPS Profile.

5 Create distributed IDS/IPS rules and publish them: Create rules to apply a previously created profile to selected applications and traffic. For more information, see Add Rules for NSX Distributed IDS/IPS and NSX Distributed Malware Prevention.

6 Verify NSX IDS/IPS status on hosts: For more information, see Verify NSX Distributed IDS/IPS Status on Host.

7 Monitor NSX IDS/IPS events. For more information, see Monitoring IDS/IPS Events.

Workflow for NSX Distributed Malware Prevention on Virtual Machine Endpoints

On the distributed east-west traffic, NSX Malware Prevention feature uses the file introspection capabilities of the NSX Guest Introspection (GI) Platform.

Important NSX Malware Prevention feature can function as designed only when your NSX environment is connected to the Internet.

To protect guest VMs on vSphere host clusters with NSX Malware Prevention feature, you must complete a series of steps.

Workflow:

1 Prepare your NSX environment for deploying the NSX Distributed Malware Prevention service. This preparation involves the following prerequisite tasks:

- Set up NSX Proxy Server for Internet Connectivity.
- Deploy NSX Application Platform.
- Activate the NSX Malware Prevention feature on the NSX Application Platform.
- Configure vSphere host clusters as NSX Host Transport Nodes by applying a Transport Node profile.
- Generate a public-private key pair for an SSH access to the NSX Malware Prevention service virtual machine. A key pair is required for logging in to the service virtual machine to download log files.
- Do a custom or a complete VMware Tools installation to install NSX File Introspection driver on VMs.
- Download the OVA file for deploying NSX Malware Prevention service virtual machine (SVM) on host clusters, which are prepared for NSX.
- Register the NSX Distributed Malware Prevention service.

For detailed instructions, see Prerequisites for Deploying the NSX Distributed Malware Prevention Service.

2 Deploy the NSX Distributed Malware Prevention service on NSX-prepared host clusters. This step turns on the NSX Malware Prevention feature on host clusters.

For detailed instructions, see Deploy the NSX Distributed Malware Prevention Service.

3 Add a security policy to protect VMs with NSX Distributed Malware Prevention service. This step involves the following Policy Management tasks:

   - Add a Malware Prevention profile.
   - Create groups and add VMs that you want to protect from malware in these groups. You can add VMs as static members, or define membership criteria that evaluate to VMs as effective members.
   - Add Distributed Malware Prevention rules. Attach the Malware Prevention profile to the rules.
   - Publish the rules to push them to the hosts.

   For detailed instructions, see Add Rules for NSX Distributed IDS/IPS and NSX Distributed Malware Prevention.

4 Monitor and analyze the file events in the NSX Manager UI.

   For detailed instructions, see Monitoring File Events.

Prerequisites for Deploying the NSX Distributed Malware Prevention Service

NSX Malware Prevention on a Distributed Firewall uses the NSX Guest Introspection (GI) framework. To detect and prevent malware on the guest endpoints (VMs), you must deploy the NSX Distributed Malware Prevention service on the ESXi host clusters that are prepared for NSX.

When you deploy the service on a host cluster, an instance of the NSX Malware Prevention service virtual machine (SVM) is deployed on each host of the cluster. Currently, an SVM of a fixed size is deployed and it requires the following resources on each host of the cluster:

- 4 vCPU
- 6 GB RAM
- 80 GB Disk space

Before deploying the NSX Distributed Malware Prevention service on host clusters, you must complete the prerequisites that are explained in the following sections. If some prerequisites are already completed, skip those, and proceed with the pending prerequisites.

Add an Appropriate License in NSX

To use the NSX Malware Prevention feature, NSX must use an appropriate license. For information about licenses that support NSX Malware Prevention, see System Requirements for NSX IDS/IPS and NSX Malware Prevention.

To add a license:

1 In NSX Manager, navigate to **System > Licenses > Add License**.

2 Enter the license key.

Verify All Hosts are Managed by VMware vCenter

NSX Malware Prevention feature is supported only on vSphere host clusters that are managed by one or multiple vCenter Servers.

1 In NSX Manager, navigate to **System > Fabric > Hosts**.

2 Ensure that you are in the **Clusters** tab.

A list of vSphere host clusters is displayed. Verify that this list includes the host clusters that are of interest to you for enabling malware protection.

Configure Hosts as Transport Nodes

Apply a Transport Node Profile to the vSphere host clusters to configure the vSphere hosts as Host Transport Nodes.

For detailed instructions, see the following topics in the *NSX Installation Guide*:

- Add a Transport Node Profile
- Prepare ESXi Hosts as Transport Nodes

Generate Public-Private Key Pair for SSH Access to SVM

To download log file from the SVM for troubleshooting purposes, read-only SSH access to the NSX Malware Prevention SVM is required.

SSH access to the **admin** user of the SVM is key-based (public-private key pair). A public key is needed when you are deploying the service on an ESXi host cluster, and a private key is needed when you want to start an SSH session to the SVM.

You can generate the public-private key pair by using any SSH key generation tool. However, the public key must adhere to a specific format, as described in the following subsection. Examples of SSH key generation tools are: ssh-keygen, PuTTY Key Generator, and so on. Supported key sizes are 1024 bits, 2048 bits, and 4096 bits.

Public Key Format

The public key must adhere to the following format:

Example:

```
ssh-rsa
A1b2C3d4E5+F6G7XxYyZzaB67896C4g5xY9+H65aBUyIZzMnJ7329y94t5c%6acD+oUT83iHTR870973TGReXpO67U=
 rsa-key-20121022
```

If you are using PuTTY Key Generator, ensure that the public key is copied directly from the UI. If the key pair exists, first load the private key file in the PuTTY Key Generator UI, and then copy the public key that is displayed in the **Key** text box. Avoid copying the contents from a public key file. The copied contents can take a different format and might not work for the SVM.

If you are generating the key pair by using ssh-keygen utility on Linux systems, the key format always includes ssh-rsa in the public key. Therefore, on Linux systems, you can copy the contents from a public key file.

Recommended Practice

NSX Distributed Malware Prevention service deployment is done at the level of a host cluster. So, a key pair is tied to a host cluster. You can create either a new public-private key pair for a service deployment on each cluster, or use a single key pair for service deployments on all the clusters.

If you plan to use a different public-private key pair for service deployment on each cluster, ensure that the key pairs are named correctly for easy identification.

A good practice is to identify each service deployment with a "compute cluster id" and specify the cluster id in the name of the key pair. For example, let us assume that the cluster id is "1234-abcd". For this cluster, you can specify the service deployment name as "MPS-1234-abcd", and name the key pair to access this service deployment as "id_rsa_1234_abcd.pem". This practice makes it easy for you to maintain and associate keys for each service deployment.

Important Store the private key securely. Loss of the private key can lead to a loss of SSH access to the NSX Malware Prevention SVM.

Deploy NSX Application Platform

NSX Application Platform is a modern microservices platform that hosts several NSX features that collect, ingest, and correlate network traffic data.

For detailed instructions about deploying the platform, see the *Deploying and Managing the VMware NSX Application Platform* publication at https://docs.vmware.com/en/VMware-NSX/index.html. From the left navigation pane at this link, expand version 4.0 or later, and then click the publication name.

Activate NSX Malware Prevention Feature

For detailed instructions, see Activate NSX Malware Prevention .

When this feature is activated, the microservices that are required for NSX Malware Prevention start running in the NSX Application Platform.

Before proceeding to the next step, verify the status of the NSX Malware Prevention feature on the NSX Application Platform. Do these steps:

1 In NSX Manager, navigate to **System > NSX Application Platform**.

2 Scroll down the page until you see the **Features** section.

3 Verify that the **NSX Malware Prevention** feature card shows **Status** as `Up`.

If the status is `Down`, wait until the status changes to `Up`, and then proceed to the next step.

Verify VM Hardware Configuration on Guest VMs

Verify that VM Hardware Configuration version 9 or later is running on the guest VMs. Do these steps:

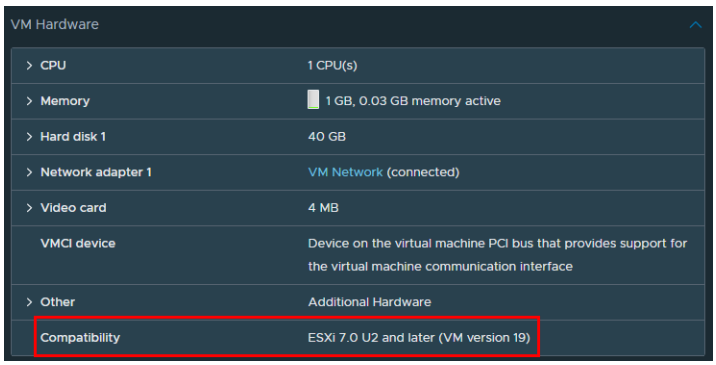
1 Log in to the vSphere Client.

2 Go to **Hosts and Clusters** and navigate to the cluster.

3 Click the VMs in the cluster, one at a time.

4 On the **Summary** page, expand the **VM Hardware** pane, and observe the Compatibility information of the VM. The VM version number must be 9 or later.

For example:



Install NSX File Introspection Driver

NSX File Introspection driver is included with VMware Tools for Windows. However, this driver is not a part of the default VMware Tools installation. To install this driver, you must do a custom or a complete installation and select the NSX File Introspection driver.

The File Introspection driver for Linux is available as part of the operating system specific packages (OSPs). The packages are hosted on VMware packages portal. Enterprise or Security Administrator (non-NSX Administrator) can install the Guest Introspection thin agent on Linux guest VMs outside of NSX. Installing open-vm-tools or VM Tools is not required for Linux.

For more information, see:

- Install the Guest Introspection Thin Agent on Windows Virtual Machines for Network Introspection.
- Install the Guest Introspection Thin Agent for Anti-virus on Linux Virtual Machines.

Download the OVA File of NSX Malware Prevention Service Virtual Machine

1 In a Web browser, open the **All Downloads** page on the VMware Customer Connect portal, and log in with your VMware ID.

2 From the **All Products** drop-down menu, select **Networking & Security**.

3 Next to VMware NSX®, click **Download Product**. The **Download VMware NSX** page opens.

4 Find the NSX license that you are using, and then click **Go to Downloads**.

5 Download the OVA file of the NSX SVM Appliance (`VMware-NSX-Malware-Prevention-appliance-version_number.build_number.ova`).

6 Extract the OVA file with the following command:

```
tar -xvf filename.ova
```

Replace *filename* with the exact name of the OVA file that you downloaded in the previous step.

Observe that the following four files are available in the root directory where the OVA file is extracted.

- OVF file (`.ovf`)
- Manifest file (`.mf`)
- Certificate file (`.cert`)
- Virtual machine disk file (`.vmdk`)

7 Copy all the extracted files to a Web server that meets the following prerequisites:

- The Web server must have unauthenticated access over HTTP.
- The Web server must be accessible to NSX Manager, all ESXi hosts where you plan to deploy the NSX Malware Prevention SVM, and the VMware vCenter that is registered to NSX.
- The MIME types for the extracted files must be added to the Web server. For information about adding MIME types to the Web server, see your Web server documentation.

File Extension	MIME Type
`.ovf`	`application/vmware`
`.vmdk`	`application/octet-stream`
`.mf`	`text/cache-manifest`
`.cert`	`application/x-x509-user-cert`

Note You can deploy the Web server on the same network where the NSX Manager appliances, ESXi hosts, and the VMware vCenter appliance are deployed. The Web server does not require Internet access.

Register the NSX Distributed Malware Prevention Service

Run the following POST API:

```
POST https://{nsx-manager-ip}/napp/api/v1/malware-prevention/svm-spec
```

In the request body of this POST API, specify the following details:

- Complete path to the OVF file on the Web server

- Name of the deployment specification (SVM is identified by this name on the VMware vCenter)
- SVM version number

Example Request Body:

```
{
    "ovf_url" : "http://{webserver-ip}/{path-to-ovf-file}/{filename}.ovf",
    "deployment_spec_name" : "NSX_Distributed_MPS",
    "svm_version" : "3.2"
}
```

The `svm_version` parameter in this POST API shows a sample value. You can specify the value of the SVM appliance version that you have downloaded.

For more information about this API including an example response, see the Malware Prevention API documentation on the VMware Developer Documentation portal.

Verify that the service name is listed on the **Catalog** page. Do these steps:

1 In NSX Manager, navigate to **System > Service Deployments > Catalog**.

2 Verify that the **VMware NSX Distributed Malware Prevention Service** is listed on the page.

Deploy the NSX Distributed Malware Prevention Service

To activate NSX Malware Prevention on vSphere host clusters, deploy the NSX Distributed Malware Prevention service on each host cluster.

When you deploy the service on a host cluster, an instance of the NSX Malware Prevention service virtual machine (SVM) is deployed on each host of the cluster. Currently, the deployed SVM has a fixed size of 4 vCPU, 6 GB RAM, and 80 GB disk space. If you add new hosts to the cluster, an instance of the SVM is deployed automatically on the new hosts.

Prerequisites

Complete the prerequisites for deploying an NSX Malware Prevention SVM. See Prerequisites for Deploying the NSX Distributed Malware Prevention Service.

Procedure

1 From your browser, log in with **admin** privileges to an NSX Manager at https://*nsx-manager-ip-address*.

2 Navigate to **System > Service Deployments > Deployment**.

3 In the **Partner Service** drop-down menu, select **VMware NSX Distributed Malware Prevention Service**, and click **Deploy**.

4 Enter the service deployment name.

5 Select the VMware vCenter that is registered as a compute manager in your NSX.

6 Select the cluster where you want to deploy the service.

7 To specify the datastore, do one of the following actions:

- Select a shared datastore as the repository for the service virtual machines.
- Select **Specified on Host**.

  The **Specified on Host** option means that you do not need to select a datastore and network on the **Deploy Service** page. Before deploying the service, you must configure Agent VM settings on each ESXi host to point to a specific datastore and network.

  To know more about configuring Agent VM settings, see the vSphere product documentation.

8 Under Networks, click **Set** and select the Management NIC (eth0) you want to use for the deployment.

a Select the network to use for the Management interface (eth0) of the SVM.

Note The selected network must have connectivity to the management network, that is, NSX Manager nodes and the components that are running on the NSX Application Platform.

If you have set the datastore as **Specified on Host**, you must set the network also as **Specified on Host**.

b Set the Network type to **DHCP** or **Static IP Pool**. If you set the network type to a Static IP Pool, select from the list of available IP pools.

Note NSX auto-assigns the control interface IP address when the SVM is deployed. For NSX Malware Prevention, the control interface IP is 169.254.1.22.

9 In the **Deployment Template** drop-down menu, select the registered deployment template.

10 (Required) Next to **Deployment Template**, click **Configure Attributes**. In the **Appliance Public Key** text box, enter or paste the public key that you created for the host cluster while completing the prerequisites, and click **Save**.

When you specify the appliance public key, you can later log in to the appliance (SVM) on each host by using the corresponding private key, and download the SVM log file for troubleshooting purposes.

11 On the **Deployment** page, click **Save** to start the deployment process.

The deployment process might take some time. While the deployment is in progress, you can watch the progress of OVF deployment and ESX Agent installation in the **Recent Tasks** pane of the vSphere Client.

12 In the NSX Manager UI, refresh the deployment status on the **Deployment** page. Wait until the status changes to `Up`.

You might have to refresh the **Deployment** page a few times to retrieve the latest status.

If the Status column shows `Down`, click the icon next to `Down`. All deployment errors are displayed. Take the required actions to fix the errors, and click **Resolve**. The status changes to `In Progress`. Wait until the status changes to `Up`.

Results

NSX Malware Prevention SVM is deployed on all the hosts of the cluster.

What to do next

Go to the **Service Instances** page. Verify that the Deployment Status and Health Status of the service instance on each host in the cluster shows `Up`.

If you need help for resolving NSX Malware Prevention service deployment issues, see Troubleshooting NSX Malware Prevention Service Virtual Machine Problems.

Log in to the NSX Malware Prevention Service Virtual Machine

By default, an **admin** user on the NSX Malware Prevention service virtual machine (SVM) does not have an SSH access to the SVM. The VMware vCenter administrator must activate SSH access to the SVM.

SSH access to the **admin** user of the SVM is key-based (public-private key pair). A public key is needed when you are deploying the service on an ESXi host cluster, and a private key is needed when you want to start an SSH session to the SVM.

Important Store the private key securely. Loss of the private key can lead to a loss of SSH access to the NSX Malware Prevention SVM.

Prerequisites

1. The public key of the NSX Malware Prevention SVM must be specified during service deployment and the key must adhere to a specific format. For information about the public key format, see Prerequisites for Deploying the NSX Distributed Malware Prevention Service.
2. VMware vCenter administrator must activate SSH access to the NSX Malware Prevention SVM by completing these steps:
   a. Log in to vSphere Client.
   b. Go to **Hosts and Clusters** and navigate to the cluster.
   c. Select the VM (service virtual machine), and then click **Launch Web Console**.
   d. Log in to the SVM as a **root** user, and run the following command to start the SSH service:

```
/etc/init.d/ssh start
```

Note On the first login, you are prompted to reset the default password of the **root** user. The default password is **`vmware`**.

Now, you can log in to the SVM as an **admin** user and use the SVM private key to start an SSH session.

Procedure

1 Ensure that the private key file is stored on your computer from where you want to start a remote SSH session to the SVM.

For example, let us assume that you had generated an RSA public-private key pair before service deployment. The private key (`id_rsa`) is saved on your Windows computer at `C:\Users\username\.ssh`.

2 On your Windows computer, open an SSH client and do these steps.

a Enter the IP address of the SVM management interface.

b Select the SVM private key file to use for authentication to the SVM.

For example, if you are using the PuTTy client, navigate to **Connection > SSHAuth**. In the **Private key file for authentication** text box, click **Browse**, and navigate to the `C:\Users\username\.ssh\id_rsa` on your Windows computer to select the private key file.

If you are using any other SSH client, consult the documentation of your SSH client for steps about specifying the private key file.

If you are using the Mac Terminal or any SSH Terminal, run the following command to start an SSH session by using the SVM private key:

```
ssh -i path_to_private_key admin@svm-management-ip
```

- Replace *path_to_private_key* with the actual path to the folder where the private key file is stored on your machine.
- Replace *svm-management-ip* with the actual IP address of the SVM management interface.

What to do next

After the debugging or troubleshooting tasks on the SVM are done, the VMware vCenter administrator must preferably deactivate SSH access to the NSX Malware Prevention SVM.

Add Rules for NSX Distributed IDS/IPS and NSX Distributed Malware Prevention

The NSX Manager UI provides a common rule table to add rules for NSX Intrusion Detection/ Prevention and NSX Malware Prevention on a Distributed Firewall.

Prerequisites

For NSX Malware Prevention:

- NSX Malware Prevention service virtual machine is deployed on vSphere host clusters that are prepared for NSX. For detailed instructions, see Deploy the NSX Distributed Malware Prevention Service.

- Add a Malware Prevention Profile .
- Create groups and add VMs that you want to protect from malware in these groups. You can add VMs as static members, or define dynamic membership criteria that evaluate to VMs as effective members. For detailed instructions, see Add a Group.

For NSX IDS/IPS:

- Add an NSX IDS/IPS Profile.
- Turn on or activate NSX IDS/IPS on the vSphere host clusters. (**Security > IDS/IPS & Malware Prevention > Settings > Shared**).

Procedure

1 From your browser, log in to an NSX Manager at https://*nsx-manager-ip-address*.

2 Navigate to **Security > IDS/IPS & Malware Prevention > Distributed Rules**.

3 Click **Add Policy** to create a section for organizing the rules.

   a Enter a name for the policy.

   b (Optional) In the policy row, click the gear icon to configure advanced policy options. These options are applicable only to NSX Distributed IDS/IPS and not to NSX Distributed Malware Prevention.

Option	Description
Stateful	A stateful firewall monitors the state of active connections and uses this information to determine which packets to allow through the firewall.
Locked	The policy can be locked to prevent multiple users from editing the same sections. When locking a section, you must include a comment. Some roles such as enterprise administrator have full access credentials, and cannot be locked out. See Role-Based Access Control.

4 Click **Add Rule** and configure the rule settings.

   a Enter a name for the rule.

   b Configure **Sources**, **Destinations**, and **Services** columns based on the traffic that requires IDS inspection. IDS supports Generic and IP Addresses Only group types for source and destination.

   These three columns are not supported for Distributed Malware Prevention firewall rules. Retain them as Any. However, you must limit the scope of the Distributed Malware Prevention rules by selecting the groups in the **Applied To** column.

   c In the **Security Profiles** column, select the profile to use for this rule.

   You can select an NSX IDS/IPS profile or an NSX Malware Prevention profile, but not both. In other words, only one security profile is supported in a rule.

d In the **Applied To** column, select any one of the options.

Option	Description
DFW	Currently, Distributed Malware Prevention rules do not support DFW in **Applied To**. Distributed IDS/IPS rules can be applied to DFW. The IDS/IPS rules get applied to workload VMs on all host clusters that are activated with NSX IDS/IPS.
Groups	The rule is applied only to the VMs that are members of the selected groups.

e In the **Mode** column, select any one of the options.

Option	Description
Detect Only	For NSX Malware Prevention service: The rule detects malicious files on the VMs, but no preventive action is taken. In other words, malicious files are downloaded on the VMs. For NSX IDS/IPS service: The rule detects intrusions against signatures and does not take any action.
Detect and Prevent	For NSX Malware Prevention service: The rule detects known malicious files on the VMs and blocks them from being downloaded on the VMs. For NSX IDS/IPS service: The rule detects intrusions against signatures and either drops or rejects the traffic depending on the signature configuration in the IDS/IPS profile or in the global signature configuration.

f (Optional) Click the gear icon to configure other rule settings. These settings are applicable only to NSX Distributed IDS/IPS and not to NSX Distributed Malware Prevention.

Option	Description
Logging	Logging is turned off by default. Logs are stored in the `/var/log/dfwpktlogs.log` file on ESXi hosts.
Direction	Refers to the direction of traffic from the point of view of the destination object. IN means that only traffic to the object is checked. OUT means that only traffic from the object is checked. In-Out, means that traffic in both directions is checked.
IP Protocol	Enforce the rule based on IPv4, IPv6, or both IPv4-IPv6.
Oversubscription	You can configure whether excess traffic should be dropped or bypassed by the IDS/IPS engine in case of oversubscription. Value entered here will overide the value set for oversubscription in the global setting.
Log Label	Log Label is stored in the firewall log when logging is enabled.

5 (Optional) Repeat step 4 to add more rules in the same policy.

6 Click **Publish**.

The rules are saved and pushed to the hosts. You can click the graph icon to view rule statistics for NSX Distributed IDS/IPS.

Note Rule statistics for NSX Distributed Malware Prevention firewall rules are not supported.

Results

When files are extracted on the endpoint VMs, file events are generated and shown on the **Malware Prevention** dashboard and the **Security Overview** dashboard. If the files are malicious, the security policy is enforced. If the files are benign, they are downloaded on the VMs.

For rules configured with IDS/IPS profile, if the system detects malicious traffic, it generates an intrusion event and shows it on the **IDS/IPS** dashboard. The system drops, rejects, or generates an alarm for the traffic based on the action that you configured in the rule.

Example

For an end-to-end example of configuring a Distributed Firewall rule for malware detection and prevention on VM endpoints, see Example: Add Rules for NSX Distributed Malware Prevention.

What to do next

Monitor and analyze file events on the **Malware Prevention** dashboard. For more information, see Monitoring File Events.

Monitor and analyze intrusion events on the **IDS/IPS** dashboard. For more information, see Monitoring IDS/IPS Events.

Example: Add Rules for NSX Distributed Malware Prevention

In this example, your objective is to create a security policy with Distributed Malware Prevention firewall rules that detects and prevents malicious Portable Executable files on Windows workload VMs that are running database servers and Web servers in your organization.

You can use the NSX Distributed Malware Prevention service in NSX to meet this objective. For this example, you will group workload VMs of database servers and Web servers by using a dynamic membership criterion based on tags.

Assumptions:

- Tags required for grouping of workload VMs are added already in the NSX inventory, as follows:
  - Tag Name: DB, Scope: Servers
  - Tag Name: WEB, Scope: Servers
- DB Tag is assigned to three database workloads (Windows VMs): VM1, VM2, and VM3.
- WEB Tag is assigned to three application workloads (Windows VMs): VM4, VM5, and VM6
- **Cloud File Analysis** option is selected in the Malware Prevention profile.

Prerequisites

NSX Malware Prevention service virtual machine is deployed on vSphere host clusters where the workload VMs are running. For detailed instructions, see Deploy the NSX Distributed Malware Prevention Service.

Procedure

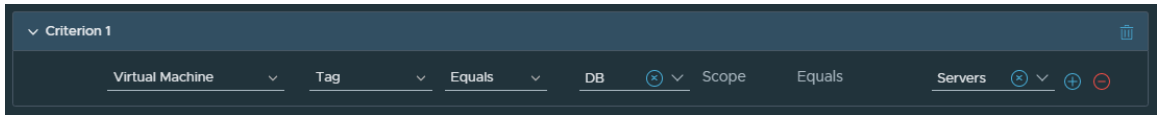
1 From your browser, log in to an NSX Manager at https://*nsx-manager-ip-address*.

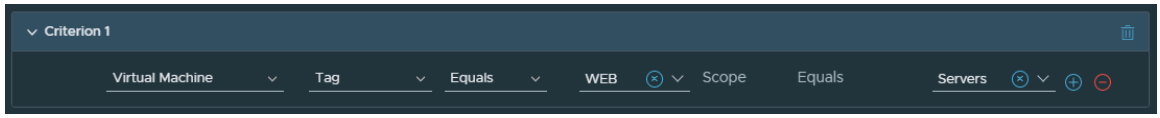
2 Organize database workload VMs and application workload VMs into two Groups.

a Navigate to **Inventory > Groups**.

b Click **Add Group**.

c Create two groups with a dynamic membership criterion based on tags and with Virtual Machine as members, as shown in the following screenshots.





d On the **Groups** page, for each group, click **View Members**, and verify that the effective members are shown.

Group Name	Effective Members
DB-Servers	VM1, VM2, VM3
Web-Servers	VM4, VM5, VM6

3 Navigate to **Security > IDS/IPS & Malware Prevention > Distributed Rules**.

4 Click **Add Policy** to create a section, and enter a name for the policy.

For example, enter **`Malware-Prevention-Rules`**.

5 Click **Add Rule** and configure the rule settings.

a Enter a name for the rule.

For example, enter **`Protect-DB-Web-Servers`**.

b In the **Sources**, **Destinations**, and **Services** columns, retain Any.

c In the **Security Profiles** column, select the Malware Prevention profile to use for this rule.

d In the **Applied To** column, select the **DB-Servers** and **Web-Servers** groups that you created earlier.

e In the **Mode** column, select **Detect and Prevent**.

6 Publish the rule.

Results

The rule is pushed to the host.

When Windows Portable Executable (PE) files are detected on the workload VMs, file events are generated and shown in the **Malware Prevention** dashboard. If the file is benign, the file is downloaded on the workload VM. If the file is a known malware (file matches known malware file signatures in NSX), and **Detect and Prevent** mode is specified in the rule, then the malicious file is blocked on the workload VM.

If the file is an unknown malware (Day-zero threat) and detected for the first time in the data center, it is downloaded on the workload VM. After NSX has determined the file verdict as malicious either by using local file analysis or cloud file analysis, the verdict is distributed to the other ESXi hosts and NSX Edges in the data center, which are activated for NSX Malware Prevention. When the file with the same hash is detected again on any of the workload VMs that are protected by NSX Malware Prevention, the security policy is applied, and the malicious file is blocked on the workload VMs.

Verify NSX Distributed IDS/IPS Status on Host

To use the NSX virtual appliance CLI, you must have SSH access to an NSX virtual appliance. Each NSX virtual appliance contains a command-line interface (CLI).

The viewable modes in the CLI can differ based on the assigned role and rights of a user. If you are unable to access an interface mode or issue a particular command, consult your NSX administrator.

Procedure

1 Open an SSH session to a compute host running the work loads that were previously deployed. Log in as root.

2 Enter the `nsxcli` command to open the NSX CLI.

3 To confirm that IDS is enabled on this host, run the command: `get ids status`.

Sample Output:

```
localhost> get ids status
 NSX IDS Status
--------------------------------------------------
 status: enabled
 uptime: 793756 (9 days 04:29:16)
```

4 To confirm both of the IDS profiles have been applied to this host, run the command `get ids profile`.

```
localhost> get ids profiles
 NSX IDS Profiles
--------------------------------------------------
Profile count: 2
 1. 31c1f26d-1f26-46db-b5ff-e6d3451efd71
 2. 65776dba-9906-4207-9eb1-8e7d7fdf3de
```

5 To review IDS profile (engine) statistics including the number of packets processed and alerts generated, run the command `get ids engine profilestats <tab_to_select_profile_ID>`.

The output is on a per profile basis, and shows the number of alerts, and the number of packets that were evaluated.

```
localhost> get ids engine profilestats eec3ea3f-0b06-4b9d-a3fe-7950d5726c7c
Fri Oct 23 2020 UTC 21:22:36.257
           NSX IDS Engine Profile Stats
------------------------------------------------------------
               Profile ID: eec3ea3f-0b06-4b9d-a3fe-7950d5726c7c
             Total Alerts: 14
            Total Packets: 27407
```

Delete the NSX Distributed Malware Prevention Service Deployment

You can turn off NSX Distributed Malware Prevention service on specific vSphere host clusters by deleting the service deployment on those host clusters.

When you delete the NSX Distributed Malware Prevention service deployment on a host cluster, an instance of the service VM on each host of that cluster is removed. The workload VMs on the hosts in that cluster lose NSX Malware Prevention security protection.

Procedure

1 From your browser, log in with **admin** privileges to an NSX Manager at https://*nsx-manager-ip-address.*

2 Navigate to **System > Service Deployments > Service Deployment**.

3 In the **Partner Service** drop-down menu, select **VMware NSX Distributed Malware Prevention Service**.

The list of service deployments is displayed.

4 Next to the service deployment name, click  .

5 In the pop-up window, click **Delete**.

The service deployment for the host cluster is permanently deleted and cannot be recovered.

If the service deployment is not deleted due to any error, the Status goes into Failed state. To completely delete a service deployment that is in a Failed state, click  , and then click **Force Delete**.

Results

NSX Malware Prevention security protection is no longer available on this host cluster. In other words, the workload VMs on this host cluster cannot detect and prevent malware.

Upgrade the NSX Malware Prevention Service Virtual Machine

You can trigger the upgrade process of the NSX Malware Prevention service virtual machine (SVM) when a new OVA file of the SVM is available on the Download VMware NSX page.

Prerequisites

- Verify that the service deployment status and the health status of the NSX Distributed Malware Prevention service instance on each host of the ESXi cluster shows `Up`.
  a. Log in to NSX Manager, and navigate to **System > Service Deployments > Deployment**.
  b. From the **Partner Service** drop-down menu, select **VMware NSX Distributed Malware Prevention Service**.
  c. Verify that the **Status** column of the service deployment in all host clusters shows `Up`.
  d. Click the **Service Instances** tab, and verify that the **Deployment Status** column and **Health Status** column of the service instance on each host of the cluster shows `Up`.
- Verify that the Kubernetes cluster on the NSX Application Platform is stable.
  a. In NSX Manager, navigate to **System > NSX Application Platform**.
  b. Verify that the cluster status is `Stable` (green).
- Verify that the status of the NSX Malware Prevention feature on the NSX Application Platform is `Up`.
  a. In NSX Manager, navigate to **System > NSX Application Platform**.
  b. Scroll down the page until you see the **Features** section.
  c. Verify that the **NSX Malware Prevention** feature card shows **Status** as `Up`.

Procedure

1. Download the new OVA file of the NSX Malware Prevention SVM.
   a. In a Web browser, open the **All Downloads** page on the VMware Customer Connect portal, and log in with your VMware ID.
   b. From the **All Products** drop-down menu, select **Networking & Security**.
   c. Next to VMware NSX®, click **Download Product**. The **Download VMware NSX** page opens.

d Find the NSX license that you are using, and then click **Go to Downloads**.

e Download the OVA file of the NSX SVM Appliance (`VMware-NSX-Malware-Prevention-appliance-version_number.build_number.ova`).

2 Extract the OVA file with the following command:

```
tar -xvf filename.ova
```

Replace *filename* with the exact name of the OVA file that you downloaded in the previous step.

Observe that the following four files are available in the root directory where the OVA file is extracted.

- OVF file (`.ovf`)
- Manifest file (`.mf`)
- Certificate file (`.cert`)
- Virtual machine disk file (`.vmdk`)

3 Copy all the extracted files to a Web server that meets the following prerequisites:

a The Web server must have unauthenticated access over HTTP.

b The Web server must be accessible to NSX Manager, all ESXi hosts where you plan to upgrade the NSX Malware Prevention SVM, and the VMware vCenter that is registered to NSX.

c The MIME types for the extracted files must be added to the Web server. For information about adding MIME types to the Web server, see your Web server documentation.

File Extension	MIME Type
`.ovf`	`application/vmware`
`.vmdk`	`application/octet-stream`
`.mf`	`text/cache-manifest`
`.cert`	`application/x-x509-user-cert`

Note You can deploy the Web server on the same network where the NSX Manager appliances, ESXi hosts, and the VMware vCenter appliance are deployed. The Web server does not require Internet access.

4 Run the following API to add a new deployment specification to an existing NSX Malware Prevention service definition:

```
POST https://{nsx-manager-ip}/napp/api/v1/malware-prevention/svm-spec
```

In the request body of this POST API, specify the following details:

- Complete path to the OVF file on the Web server

- Name of the deployment specification (SVM is identified by this name on the VMware vCenter)
- SVM version number

Example Request Body:

```
{
    "ovf_url" : "http://{webserver-ip}/{path-to-ovf-file}/{filename}.ovf",
    "deployment_spec_name" : "NSX_Distributed_MPS_2",
    "svm_version" : "3.3"
}
```

Specify a deployment specification name that is easy for you to identify when you upgrade the service appliance (SVM) in the next step. The SVM version in this request body is only an example. You must replace it with the appropriate version number.

For more information about this API including an example response, see the Malware Prevention API documentation on the VMware Developer Documentation portal.

5 Upgrade the NSX Malware Prevention service virtual machine (appliance) in the service deployment of each ESXi host cluster.

a In NSX Manager, navigate to **System > Service Deployments > Deployment**.

b From the **Partner Service** drop-down menu, select **VMware NSX Distributed Malware Prevention Service**.

c Next to the service deployment of a host cluster, click  , and then click **Change Appliance**.

Verify that the **Change Appliance** window lists the deployment specification name that you specified in the request body of the POST API.

d Select the deployment specification name.

In this case, select **NSX_Distributed_MPS_2**.

e Click **Update**.

f Repeat steps c, d, and e to upgrade the service virtual machine (appliance) in the service deployments of other ESXi host clusters.

What to do next

After the upgrade is finished, verify the connectivity status, solution status, and health status of each service instance on the ESXi host.

1 Navigate to **System > Service Deployments > Service Instances**.

2 For each service instance, verify that the **Health Status** column shows `Up`.

3 Click the icon in the **Health Status** column and verify that the following statuses are `Up`:

- Solution status
- Connectivity status between NSX Guest Introspection agent and NSX Ops agent.

For help on troubleshooting issues, see Troubleshooting NSX Malware Prevention Service Virtual Machine Problems.

Using NSX IDS/IPS and NSX Malware Prevention on a Gateway Firewall

You can use the NSX IDS/IPS feature to detect malicious traffic patterns in the north-south traffic, and use the NSX Malware Prevention feature to detect malicious files in the north-south traffic.

Workflow for NSX IDS/IPS on a Gateway Firewall

Perform the following steps to use NSX IDS/IPS on a Gateway Firewall.

1 Set up NSX Proxy Server for Internet Connectivity. NSX IDS/IPS can work in a network without Internet connectivity, but you will need to manually update the IDS/IPS signatures. For more information, see Preparing the Data Center for NSX IDS/IPS and NSX Malware Prevention.

2 Download latest signature set and configure signature settings: Download the latest signature set if you have not selected automatic download option and configure actions for signatures. For more information, see Preparing the Data Center for NSX IDS/IPS and NSX Malware Prevention.

3 Enable nodes for IDS/IPS: Select gateways on which you want to enable IDS/IPS. For more information, see Preparing the Data Center for NSX IDS/IPS and NSX Malware Prevention

> **Note** NSX IDS/IPS for a Gateway Firewall is supported only for tier-1 gateways.

4 Create IDS/IPS profiles: Create profiles to group signatures. For more information, see Add an NSX IDS/IPS Profile.

5 Create gateway IDS/IPS rules and publish them: Create rules to apply a previously created profile to selected applications and traffic. For more information, see Add Rules for NSX IDS/IPS and NSX Malware Prevention on a Gateway Firewall.

6 Monitor events on nodes. For more information, see Monitoring IDS/IPS Events.

Workflow for NSX Malware Prevention on a Gateway Firewall

To protect north-south traffic that is passing through the Gateway Firewall with the NSX Malware Prevention feature, you must complete a series of steps.

> **Important** NSX Malware Prevention feature can function as designed only when your NSX environment is connected to the Internet.

Detection of malware is supported on tier-1 gateways, but not on tier-0 gateways. Prevention of malware on the Gateway Firewall is currently not supported.

Workflow:

1 Prepare your NSX environment for NSX Malware Prevention on the Gateway Firewall. This preparation involves the following tasks:

   - Set up NSX Proxy Server for Internet Connectivity.
   - Deploy NSX Application Platform.
   - Activate the NSX Malware Prevention feature on the NSX Application Platform.
   - Turn on or activate NSX Malware Prevention on the tier-1 gateways.

   You can complete these preparation tasks by using either the **IDS/IPS & Malware Prevention Setup** wizard or the IDS/IPS & Malware Prevention **Settings** page. For more information about using the setup wizard, see Preparing the Data Center for NSX IDS/IPS and NSX Malware Prevention.

2 Add a security policy to protect traffic passing through the tier-1 gateways. This step involves the following Policy Management tasks:

   - Add a Malware Prevention profile.
   - Create groups to use them in the sources and destinations of the Gateway Firewall rules. You can add static memberships in the groups or define membership criteria.
   - Add Gateway Firewall rules on the tier-1 gateways. Attach the Malware Prevention profile to the rules.
   - Publish the rules.

   For detailed instructions, see Add Rules for NSX IDS/IPS and NSX Malware Prevention on a Gateway Firewall.

3 Monitor and analyze the file events in the NSX Manager UI.

   For detailed instructions, see Monitoring File Events.

Add Rules for NSX IDS/IPS and NSX Malware Prevention on a Gateway Firewall

The NSX Manager UI provides a common rule table to add rules for NSX Intrusion Detection/ Prevention and NSX Malware Prevention on a Gateway Firewall.

The security profiles that you add to the rule determine whether the Gateway Firewall rule enforces only NSX IDS/IPS, or only NSX Malware Prevention, or both.

Prerequisites

For NSX Malware Prevention:

- Add a Malware Prevention Profile .
- Turn on or activate NSX Malware Prevention on the tier-1 gateways. (**Security > IDS/IPS & Malware Prevention > Settings > Shared**)

For NSX IDS/IPS:

- Add an NSX IDS/IPS Profile.
- Turn on or activate NSX IDS/IPS on tier-1 gateways. (**Security > IDS/IPS & Malware Prevention > Settings > Shared**)

Procedure

1 From your browser, log in to an NSX Manager at https://*nsx-manager-ip-address*.

2 Navigate to **Security > IDS/IPS & Malware Prevention > Gateway Rules**.

3 If you want to add a policy for a specific gateway, ensure that you are in the **Gateway Specific Rules** tab, and select a gateway. If you want to add a policy for multiple gateways, ensure that you are in the **All Shared Rules** tab.

4 Click **Add Policy** to create a section for organizing the rules.

   a Enter a name for the policy.

   b (Optional) In the policy row, click the gear icon to configure advanced policy options. These options are applicable only to NSX IDS/IPS and not to NSX Malware Prevention.

Option	Description
Stateful	A stateful firewall monitors the state of active connections, and uses this information to determine which packets to allow through the firewall.
Locked	The policy can be locked to prevent multiple users from editing the same sections. When locking a section, you must include a comment.

   c Click Publish to publish the policy.

5 Click **Add Rule** and configure the rule settings.

   a Enter a name for the rule.

   b In the **Sources** column, click the edit icon, and select the groups to use as the source of the rule. If source is not specified, it defaults to Any.

   For information about adding groups, see Add a Group.

   c In the **Destinations** column, click the edit icon, and select the groups to use as the destination of the rule. If destination is not specified, it defaults to Any.

d In the **Services** column, click the edit icon, and select the services to use in the rule. If service is not specified, it defaults to Any.

Note

- On clicking the edit icon, the UI displays a list of all available services. However, NSX Malware Prevention currently supports detection of file transfer only for the following services: HTTP, HTTPS, FTP, and SMB.
- NSX Malware Prevention on the Gateway Firewall currently does not support extracting and analyzing files that are uploaded using HTTP. However, if files are uploaded using FTP, the extraction and analysis of the files for detecting malicious behavior is supported.

e In the **Security Profiles** column, click the edit icon, and select the profiles to add to the firewall rule.

You can select a maximum of two security profiles–one NSX IDS/IPS profile and one NSX Malware Prevention profile.

f If you are adding the rule for a specific gateway, the **Applied To** column displays the name of that gateway.

If you are adding shared rules, click the edit icon in the **Applied To** column, and select the gateways to which you want to apply the rule.

By default, gateway firewall rules are applied to all the available uplink interfaces and service interfaces on the selected gateways.

g In the **Mode** column, select any one of the options.

Option	Description
Detect Only	The rule detects malicious files, malicious traffic, or both, on the selected gateways depending on the profile that is attached to the rule. No preventive action is taken.
Detect and Prevent	NSX Malware Prevention currently does not support this mode. However, rules with NSX IDS/IPS profile can detect and block malicious traffic on the selected gateways.

h (Optional) Click the gear icon to configure other rule settings. These settings are applicable only to NSX IDS/IPS and not to NSX Malware Prevention.

Option	Description
Logging	Logging is turned off by default. Logs are stored in the `/var/log/dfwpktlogs.log` file on ESXi hosts.
Direction	Refers to the direction of traffic from the point of view of the destination object. IN means that only traffic to the object is checked. OUT means that only traffic from the object is checked. In-Out, means that traffic in both directions is checked.
Oversubscription	Configure whether excess traffic should be dropped or should bypass the IDS/IPS engine in case of oversubscription. Value entered here will overide the value set for oversubscription in the global setting.
IP Protocol	Enforce the rule based on IPv4, IPv6, or both IPv4-IPv6.

6 (Optional) Repeat step 4 to add more rules in the same policy.

7 Click **Publish**. You can click the graph icon to view rule statistics for NSX IDS/IPS on Gateway Firewall.

The rules are saved and pushed to the NSX Edges.

Results

When files are detected on the tier-1 gateways, file events are generated and shown on the **Malware Prevention** dashboard and the **Security Overview** dashboard.

For rules configured with IDS/IPS profile, if the system detects malicious traffic, it generates an intrusion event. You can view the event details on the **IDS/IPS** dashboard or the **Security Overview** dashboard.

Example

For an end-to-end example of configuring Gateway Firewall rules with NSX Malware Prevention, see Example: Add Rules for NSX Malware Prevention on a Gateway Firewall.

What to do next

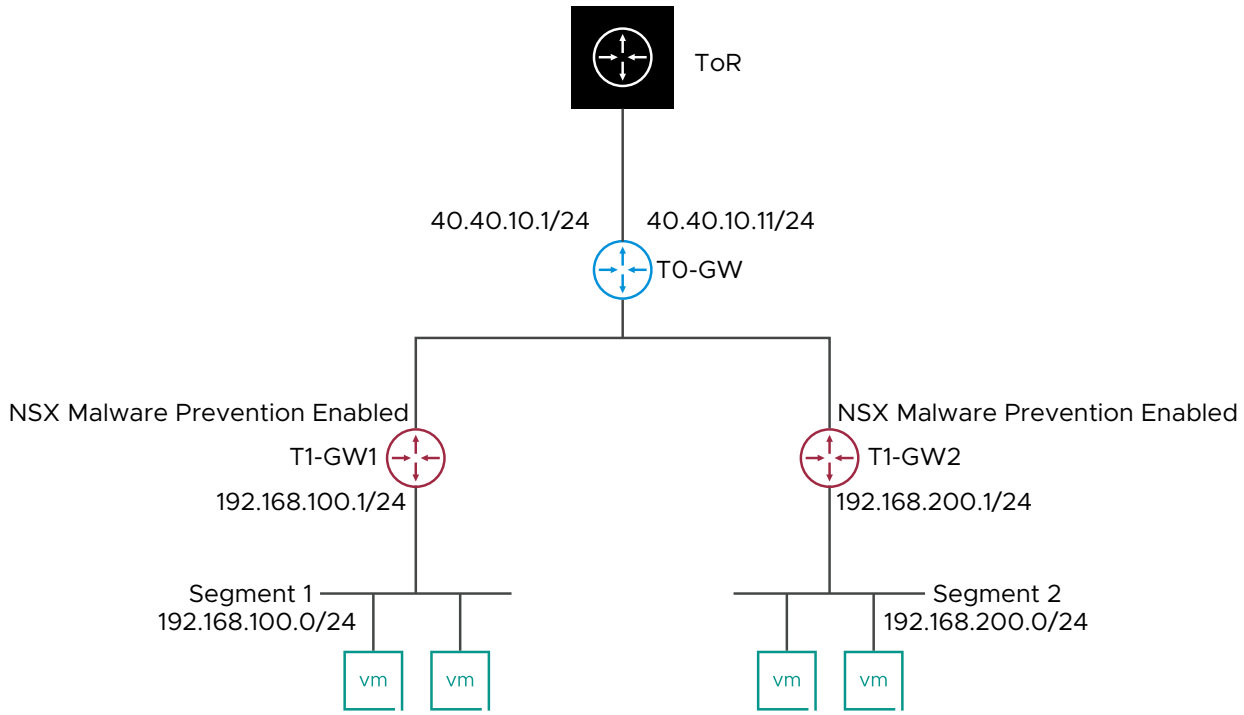
Monitor and analyze file events on the **Malware Prevention** dashboard. For more information, see Monitoring File Events.

Monitor and analyze intrusion events on the **IDS/IPS** dashboard. For more information, see Monitoring IDS/IPS Events.

Example: Add Rules for NSX Malware Prevention on a Gateway Firewall

In this example, your objective is to create security policies with Gateway Firewall rules that detect malicious files on the north-south traffic, which is passing through the NSX Edges in your NSX.

For this example, consider that your network topology is as shown in the following figure. You will add Gateway Malware Prevention rules to detect malware on tier-1 gateways: T1-GW1 and T1-GW2. Both tier-1 gateways have an overlay segment attached to it. Workload VMs are attached to the overlay segments. Both tier-1 gateways are connected to a single tier-0 gateway, which in turn is connected to the physical top-of-rack switch to enable connectivity with the outside public network.



Assumptions:

- The following groups are added in the NSX inventory.

Group Name	Group Type	Notes
North	IP Addresses Only	This group contains a public IP range. For example, 12.1.1.10-12.1.1.100
South	Generic	This group contains an overlay segment (Segment1), which is attached to T1-GW1, as the static member.

- A Malware Prevention profile named **`Profile_T1-GW`** is added with the following configuration:
  - All file category options are selected.
  - **Cloud File Analysis** option is selected.

  You will use this Malware Prevention profile in the Gateway Firewall rules of both tier-1 gateways.

Prerequisites

- NSX Edges with Extra Large form factor are deployed in your data center and configured as Edge Transport Nodes.
- NSX Malware Prevention feature is turned on or activated on tier-1 gateways: T1-GW1 and T1-GW2.

Procedure

1. From your browser, log in to an NSX Manager at https://*nsx-manager-ip-address*.
2. Navigate to **Security > IDS/IPS & Malware Prevention > Gateway Rules**.
3. On the **Gateway Specific Rules** page, in the **Gateway** drop-down menu, select **T1-GW1**.
4. Click **Add Policy** to create a section, and enter a name for the policy.

   For example, enter **`Policy_T1-GW1`**.
5. Click **Add Rule** and configure two rules with the following configurations.

Name	ID	Sources	Destinations	Services	Security Profiles	Applied To	Mode
N_to_S	1011	North	South	HTTP	Profile_T1-GW	T1-GW1	Detect Only
S_to_N	1010	South	North	HTTP	Profile_T1-GW	T1-GW1	Detect Only

   The rule IDs in this table are only for reference. They might vary in your NSX environment.

   Let us understand the meaning of these rules:

   - Rule 1011: This rule is enforced on T1-GW1 when HTTP connections are initiated by the machines in the public IP range (12.1.1.10-12.1.1.100) and these connections are accepted by the workload VMs that are attached to Segment1. If a file is detected in the HTTP connection, a file event is generated, and the file is analyzed for malicious behavior.
   - Rule 1010: This rule is enforced on T1-GW1 when HTTP connections are initiated by the workload VMs on Segment1 and these connections are accepted by the machines in the public IP range (12.1.1.10-12.1.1.100). If a file is detected in the HTTP traffic, a file event is generated, and the file is analyzed for malicious behavior.
6. Publish the rules.

7 On the **Gateway Specific Rules** page, in the **Gateway** drop-down menu, select **T1-GW2**.

8 Click **Add Policy** to create a section, and enter a name for the policy.

For example, enter **`Policy_T1-GW2`**.

9 Click **Add Rule** and configure an Any-Any rule as follows.

Name	ID	Sources	Destinations	Services	Security Profiles	Applied To	Mode
Any_Traffic	1006	Any	Any	Any	Profile_T1-GW	T1-GW2	Detect Only

This rule is enforced on T1-GW2 when any type of traffic is initiated from any source and accepted by any destination. If a file is detected in the traffic, a file event is generated, and the file is analyzed for malicious behavior.

10 Publish the rules.

Example

Scenario: In the same topology as shown earlier, assume that a VM on Segment1 wants to transmit a file to a VM on Segment2. In this case, the file traverses through both tier-1 gateways: T1-GW1 and T1-GW2. As Malware Prevention profile is configured on both tier-1 gateways, the file is inspected twice and two file events are generated. This behavior is expected.

Distributed IDS/IPS Logs

When logging is enabled for NSX-T IDS/IPS, you can look at log files to troubleshoot issues.

Below is a sample log file for NSX-T IDS/IPS, located in /var/log/nsx-idps/nsx-idps-events.log:


```
{"timestamp":"2021-08-10T01:01:15.431231+0000","flow_id":1906423505866276,"pcap_cnt":40,"event
_type":"alert","src_ip":"192.
168.100.166","src_port":49320,"dest_ip":"185.244.30.17","dest_port":1965,"proto":"TCP","direct
ion":"to_server","metadata":
{"flowbits":
["LL.verifier_tcp_successful","LL.verifier_tcp_failed","LL.verifier_tcp_blocked"],"flowints":
{"intraflow_beacon_num_strides":0,"intraflow_beacon_last_ts":1628557275,"intraflow_beacon_pack
ets_seen":1,"intraflow_beacon_grp_1"
:1,"intraflow_beacon_grp_1_cnt":0,"intraflow_beacon_grp_2":1,"intraflow_beacon_grp_2_cnt":0,"i
ntraflow_beacon_grp_3":1,
"intraflow_beacon_grp_3_cnt":0,"intraflow_beacon_prior_seq":1762155507,"intraflow_beacon_prior
_ack":1700774517,
"intraflow_beacon_num_runts":0,"intraflow_beacon_sni_seen":0}},"nsx_metadata":
{"flow_src_ip":"192.168.100.166",
"flow_dest_ip":"185.244.30.17","flow_dir":2,"rule_id":1001,"profile_id":"f7169d04-81bf-4c73-94
66-b9daec6220de",
"user_id":0,"vm_uuid":"b1396a3e-3bf9-4fd7-839d-0709c86707b0"},"alert":
{"action":"allowed","gid":1,"signature_id":1096797,"rev":14556,
"signature":"LASTLINE Command&Control: (RAT) Remcos RAT","category":"A Network Trojan was
Detected","severity":1,"source":{"ip":"185.244.30.17","port":1965},"target":
{"ip":"192.168.100.166","port":49320},
"metadata":{"detector_id":["96797"],"severity":["100"],"confidence":["80"],"exploited":
```

```
["None"],"blacklist_mode":["REAL"],"ids_mode":["REAL"],"threat_name":["Remcos RAT"],
"threat_class_name":["command&control"],"server_side":["False"],"flip_endpoints":
["False"],"ll_expected_verifier":["default"]}},
"flow":
{"pkts_toserver":3,"pkts_toclient":1,"bytes_toserver":808,"bytes_toclient":66,"start":"2021-08
-10T01:01:15.183844+0000"}}
```


Field	Description
Timestamp	The timestamp of the packet on top of which the alert was triggered.
flow_id	The unique identifier for each flow tracked by nsx-idps.
event_type	The type of event generated by the IDPS engine. For alerts, the event type will always be "alert" (regardless of the action performed).
src_ip	The source IP of the packet on top of which the alert triggered. Depending on the alert characteristics, this might be the address of the client, or the address of the server. Refer to the field "direction" to determine the client.
src_port	The source port of the packet on top of which the alert triggered.
dest_ip	The destination IP of the packet on top of which the alert triggered.
dest_port	The destination port of the packet on top of which the alert triggered.
proto	The IP protocol of the packet on top of which the alert triggered.
direction	The direction of the packet compared to the flow direction. The value will be "to_server" for a packet flowing from client to server, and "to_client" for a packet flowing from server to client.

Any fields not included on the NSX Metadat table are for internal use only.

NSX Metadata	Description
metadata.flowbits and metadata.flowints	This field constitutes a dump of the internal flow state. The variables are dynamically set by various signatures or Lua scripts operating on the specific flow. The semantics and nature of the fields are primarily internal, and may vary across IDS bundles updates.
nsx_metadata.flow_src_ip	The IP address of the client. Can be derived by looking at the packets endpoints, and at the packet direction.
nsx_metadata.flow_dest_ip	The IP address of the server.
nsx_metadata.flow_dir	The direction of the flow with respect to the originating virtual machine. Value is 1 for flows that are inbound to the monitored virtual machine, and 2 for flows that are outbound to the monitored virtual machine.
nsx_metadata.rule_id	The DFW::IDS rule ID to which the packet matched.

NSX Metadata	Description
nsx_metadata.profile_id	The context profile ID that was used by the matched rule.
nsx_metadata.user_id	The user ID whose traffic generated the event.
nsx_metadata.vm_uuid	The identifier of the virtual machine whose traffic generated the event.
alert.action	The action performed by nsx-idps on packet (Allowed/ Blocked). Depends on the configured Rule Action.
alert.gid, alert.signature_id, alert.rev	The identifier of the signature, and its revision. A signature can maintain the same identifier, and be updated to a newer version by increasing the revision.
alert.signature	A short description of the detected threat.
alert.category	The category of the detected threat. This is usually a very coarse/inaccurate categorization. Mode details can be found in alert.metadata.
alert.severity	The priority of the signature, as derived from the alert category. Higher priority alerts are usually associated with more severe threats.
alert.source/alert.target	Information on the attack direction, which is not necessarily matching the flow direction. The source of the alert will be the attacking endpoint, while the target of the alert will be the victim of the attack.
alert.metadata.detector_id	An internal identifier of the detection used by the NDR component to associate threat metadata and documentation.
alert.metadata.severity	0-100 range of the severity of the threat. This value is a function of the alert.metadata.threat_class_name.
alert.metadata.confidence	0-100 range of the degree of confidence in the correctness of the detection. Signatures that are released despite the potential for false positives report a low degree of confidence (<50).
alert.metadata.exploited	A modifier to express whether the attacker reported in the detection is likely to be a compromised host (i.e. endpoint information should not be considered a reliable IoC).
alert.metadata.blacklist_mode	Internal only.
alert.metadata.ids_mode	The operation mode for the signature. Current possible values are REAL (produces real-mode detections in the NDR product), and INFO (produces info-mode detections in the NDR product).
alert.metadata.threat_name	The name of the detected threat. The threat name is curated in the context of the NDR product as part of a well defined ontology, and is the most reliable source of information on the nature of the attack.
alert.metadata.threat_class_name	Name of the high level class of the attack to which the threat pertains. Threat classes are high level categories with values such as "command&control", "drive-by," and "exploit."

NSX Metadata	Description
alert.metadata.server_side	A modifier to express whether the threat is meant to effect servers or clients. It is equivalent to the information expressed by the alert.source, and alert.target attributes.
alert.metadata.flip_endpoints	A modifier to express whether the signature is expected to match on packets flowing from server to client, rather than client to server.
alert.metadata.ll_expected_verifier	Internal only.
flow.pkts_toserver/flow.pkts_toclient/flow.bytes_toserver/ flow.bytes_toclient	Information on the number of packets/bytes that were seen in a given flow at the time of the alert. Note that this information does not express the total amount of packets belonging to the flow. This information expresses the partial counts at the moment in which the alert was generated.
flow.start	The timestamp of the first packet belonging to the flow.

Monitoring File Events

File events are generated when files are extracted by the IDS engine on the NSX Edges in the north-south traffic and by the NSX Guest Introspection agent on the virtual machine endpoints in the distributed east-west traffic.

NSX Malware Prevention feature inspects the extracted files to determine whether they are benign, malicious, or suspicious. Each unique inspection of a file is counted as a single file event in NSX. In other words, a file event refers to a unique file inspection. The terms "file event" and "file inspection" are used interchangeably throughout this documentation.

This documentation explains the file event monitoring tasks by using the NSX Manager UI. For documentation on doing this task by using the NSX Malware Prevention file event APIs, see the VMware Developer Documentation portal.

Monitor Statistics of File Events on the Security Overview Dashboard

Use the **Security Overview** dashboard to view the summary-level statistics of file inspections in the NSX.

You can filter the file inspection statistics (file events statistics) on the dashboard for a specific time period. The default time period for each graph on the dashboard is last one hour. Maximum supported time period for each graph is last 14 days.

The following file events statistics for the selected time period are shown in a graphical format on this dashboard:

- Total number of inspected file events, malicious file events, suspicious file events, and blocked files.
- Number of file inspections for different ranges of threat score.
- Top five recently inspected files in the data center sorted by the timestamp.
- Top five malicious files detected in the data center.

- Trend of malicious file events, suspicious file events, and suppressed file events in the data center.
- Distribution of file inspections based on the malware family to which the files belong.
- Breakdown of file inspections by the type of analysis performed (local file analysis, cloud file analysis).

Prerequisites

- NSX Malware Prevention feature is activated successfully in the NSX Application Platform.
- NSX Malware Prevention feature is activated on the ESXi host clusters or tier-1 gateways, or both, depending on your security requirements.

Procedure

1. From your browser, log in to an NSX Manager at https://*nsx-manager-ip-address*.
2. Navigate to **Security > Security Overview > Threat Detection & Response > Malware Prevention**.

   File events statistics are displayed. By default, all the graphs show statistics for the last one week.
3. At the top-right corner of each graph, click the drop-down menu to filter the statistics for the time period that you are interested in.
4. Point to the various data points in the graphs to view additional information as tooltips.
5. (Optional) Click the linked texts on this dashboard to jump to the other dashboard pages and drill down to the inspection details and history of inspections for specific files.

   For example:

   - Click the **Malicious File Events** hyperlinked text to jump to the **Potential Malware** page on the **Malware Prevention** dashboard.
   - Click the **Inspected File Events** hyperlinked text to jump to the **All Files** page on the **Malware Prevention** dashboard.

Monitor Details of File Events on the Malware Prevention Dashboard

Use the **Malware Prevention** dashboard to drill down to events details of files that are extracted in the data center for deeper monitoring and analysis purposes.

The dashboard can show file events over the last 14 days. For information about the maximum number of file events that are supported on the Distributed Firewall and Gateway Firewall, see the VMware Configuration Maximums tool at https://configmax.vmware.com/home.

The information about file events (file inspections) is shown in two tab pages.

Potential Malware page

Shows aggregated events details of malicious files, suspicious files, and allowlisted files that are extracted in the data center over a specific time period.

A bubble in the bubble chart represents a unique file that is extracted in the data center. A file is uniquely identified by its file hash. The color and the graphic inside the bubble denote whether the file is malicious, suspicious, or allowlisted.

A row in the table represents one file. The number on the bubble denotes the threat score computed for the file. The score ranges from 0–100, and it denotes the degree of risk or malicious intent that is associated with the file. A high threat score indicates a greater amount of risk, and the reverse. For example:

- Score range for benign files is 0–29.
- Score range for suspicious files is 30–69.
- Score range for malicious files is 70–100.
- Uninspected files have a score of -1.

If the verdict of the file is malicious or suspicious, the malware family and malware class for that file is displayed. A single file can belong to multiple malware families and malware classes. However, if malware family and malware class for a file is unknown to NSX, the information is not displayed in the UI.

Note For each file, the event details (inspection details) are aggregated and shown on the dashboard. For example, if a single file is inspected five times in the data center, five file events are generated. In other words, the count of inspections for the file is five. However, the bubble chart shows a single bubble for the file, and the table has a single row for that file. When you point to a bubble, a summary of inspections done for the file is shown. Similarly, when you expand the row for a file in the table, the details of the most recent file inspection are shown. Nevertheless, the history of all previous inspections for the file is retained and available for you to see.

The following table describes the meaning of the icons used on the bubble chart.

Icon	Meaning
	A small bubble on the timeline represents a single inspection for a file.
	A large bubble on the timeline represents multiple inspections for a single file. Example: Assume that an `.exe` file is extracted on five guest VMs over three days, and NSX has determined this file as suspicious. In this case, five unique file inspections have occurred for the `.exe` file in the data center. A large bubble is shown on the suspicious timeline on the last inspected timestamp. You can click the bubble to view the history of all five inspections for this `.exe` file.
	A group of bubbles on the timeline represents multiple unique file inspections with the same verdict. Example: Assume that four unique `.docx` files A, B, C, and D are extracted from the north-south traffic in the data center at the same time (or nearly the same time), and NSX has determined that all these files are malicious. The bubbles for all the four files are grouped together and shown on the malicious timeline of the bubble chart.

All Files page

Shows a tabular view of all the unique files that are extracted in the data center, including the benign files. In other words, this page shows all the unique files regardless of the verdict of the file. Expand a row in the table to view the last inspection details of the file.

Prerequisites

- NSX Malware Prevention feature is activated successfully in the NSX Application Platform.
- NSX Malware Prevention feature is activated on the ESXi host clusters or tier-1 gateways, or both, depending on your security requirements.

Procedure

1. From your browser, log in to an NSX Manager at https://*nsx-manager-ip-address*.
2. Click **Security**, and then in the left navigation pane, click **Malware Prevention**.

   The **Potential Malware** page is displayed. By default, the bubble chart and the table show files that are extracted in the last one hour. To view the files for a different time period, click the drop-down menu at the top-right corner of this page, and select a different time period.
3. (Optional) Click the filter icon at the top-right corner of the page, and select the criteria to filter the information on the page.

   The filter criteria are applied to both the bubble chart and the table. The following filter criteria are supported:

   - Verdict (including Allowlist)
   - SHA256 hash

- Blocked
- File Type
- Malware Class
- Malware Family

4 Monitor the file events details (inspections) that are shown on the dashboard.

a Point to a bubble to view the summary information about the inspections for a file in a pop-up window.

The information in the pop-up window varies depending on whether you point to a small bubble, a large bubble, or a group of bubbles. For example, when you point to a small bubble, the pop-up window displays summary information about a single inspection of the file.

b Drag the timeline in the bubble chart to zoom out or zoom in, if required.

c Click a bubble to jump directly to that file in the table. Expand the row to view complete details about the most recent inspection for this file.

Field	Description
File Type	The type of file that is extracted on the transport node (host or edge). For example, PdfDocFile, PeExeFile, ShellScriptFile, and so on.
File Type Details	Brief information about the file type.
Client (Last)	The destination machine that received the file in the last inspection. For files that are extracted on the endpoint VMs in the distributed east-west traffic within the data center, the client is the endpoint VM itself. For files that are extracted on the NSX Edges in the north-south traffic, the direction of traffic determines the client. For example, if a VM inside the data center is uploading a file to a machine outside the data center, the client is the machine outside the data center. If a VM inside the data center is downloading a file from a machine outside the data center, the client is the VM inside the data center.
Server (Last)	The source machine from where the file was received in the last inspection. For files that are extracted on the endpoint VMs in the distributed east-west traffic within the data center, NSX Malware Prevention cannot determine the source of the file. Therefore, the Server (Last) box is always empty. For files that are extracted on the NSX Edges in the north-south traffic, the direction of traffic determines the server. For example, if a VM inside the data center is downloading a file from a machine outside the data center, the server is the machine outside the data center. If a VM inside the data center is uploading a file to a machine outside the data center, the server is the VM inside the data center.
File Name	The names associated with the file. A single file has a unique hash, but the clients that received the file might save the file with different names.
Protocol	The protocol used for the file transfer. For example, HTTP, FTP, HTTPS, and so on.
Workloads	Click the number next to this field to view the list of all workload VMs in the data center that are affected by the file.
Total Inspections	Click the number next to this field to view the history of all inspections done for the file. For example, if the file is inspected 10 times in the data center, the pop-up window shows a summary of all 10 inspections.
Firewall Type	The value is either `Host` or `Edge`. If the file was last extracted from the ESXi host where the Distributed Firewall is running, the value is `Host`. If the file was last extracted from the edge where the Gateway Firewall is running, the value is `Edge`.

Field	Description
Transport Node	The ID of the Edge Transport Node or the Host Transport Node where the file was extracted in the last inspection.
First Inspected	The date and time when the file was first inspected in the data center.
Last Inspected	The date and time when the file was last inspected in the data center.
Submitted By	The value is always `System`, which means that NSX has submitted the file to the cloud for a detailed analysis.
Analyst UUID	The UUID of the file submission to the cloud for a detailed analysis. The UUID is displayed regardless of whether the file is submitted to the cloud either during the last inspection or in any of the previous inspections. If the file was submitted to the cloud multiple times, the UUID of the last submission is displayed.
Blocked	Denotes whether the file is blocked. Value is either Yes or No.

d (Optional) Perform the following additional tasks:

- Add a File to the Allowlist
- View File Analysis Report
- View File Event Details in NSX Network Detection and Response UI
- View Campaign Details in NSX Network Detection and Response UI

5 Click the **All Files** tab.

This page shows a list of all the unique files that are extracted in the data center regardless of the verdict of the file. By default, files extracted in the last one hour are shown. To view the list of files for a different time period, click the drop-down menu at the top-right corner of this page, and select a different time period.

Add a File to the Allowlist

You can override or suppress the file verdict that NSX has computed. Files with suppressed verdict are listed in the Allowlist table.

Assume that NSX Malware Prevention has extracted an executable file on the guest VMs of some users in the data center, and the verdict for this file is computed as malicious. If the rule in your security policy is set to Detect and Prevent mode, NSX Malware Prevention feature blocks this file on the guest VMs. However, if you have determined that the file is legitimate and not harmful to the users in the data center, you can suppress (override) the NSX computed verdict. The allowlisted verdict is logged in the NSX Manager database for auditing purposes. When this allowlisted file is detected or extracted again in the data center, NSX Malware Prevention does not analyze this file, and returns the verdict as `Uninspected`. NSX Malware Prevention can analyze this suppressed file again in subsequent file extractions only after you remove the file from the allowlist table.

Procedure

1 From your browser, log in to an NSX Manager at https://*nsx-manager-ip-address*.

2 Click **Security**, and then in the left navigation pane, click **Malware Prevention**.

The **Potential Malware** page is displayed. You can do the subsequent steps on this page or on the **All Files** page.

3 Click the filter icon at the top-right corner of the page, and select the criteria to filter the information on the page.

Filtering the information can help you quickly find the file that is of interest to you. For example, you can select the **Verdict** criterion, and then select the **Malicious** option to view files with only malicious verdict on the page.

4 In the table, click the **Suppress** icon for the file whose verdict you want to suppress, and then click **Apply**.

Note You can suppress only one file at a time. A batch suppression of multiple files simultaneously is not supported currently.

A new file event is generated for the same file hash with verdict as **Allowlist**. The threat score of the file does not change. However, the color of the bubble changes to gray and the bubble of this file hash moves to the Allowlist timeline in the bubble chart.

5 Verify that the suppressed file is added to the Allowlist table.

a Navigate to **Security > IDS/IPS & Malware Prevention > Settings > Malware Prevention**.

b Click the **Refresh** icon at the bottom of the table.

The file is shown in the Allowlist table.

6 If you want to remove the file from the Allowlist table, select the file, and click **Delete**. In the information message that appears, click **Yes** to confirm the delete action.

View File Analysis Report

Cloud file analysis produces a report that contains detailed results of an analysis submission.

You can view this analysis report from the **Potential Malware** page or the **All Files** page on the **Malware Prevention** dashboard.

Procedure

1 From your browser, log in to an NSX Manager at https://*nsx-manager-ip-address*.

2 Click **Security**, and then in the left navigation pane, click **Malware Prevention**.

The **Potential Malware** page is displayed. You can do the subsequent steps either on this page or on the **All Files** page.

3 (Optional) Click the filter icon at the top-right corner of the page, and select the criteria to filter the information displayed on the page.

   Filtering the information can help you quickly find the file that is of interest to you. For example, you can select the **Verdict** criterion, and then select the **Malicious** option to view files with only malicious verdict on the page.

4 In the table, expand the row of the file.

   When a file hash is submitted to the cloud for a detailed analysis either during the last inspection or in any of the previous inspections, the **Analyst UUID** field displays a value. The presence of this value indicates that a file analysis report is available for this file.

5 Click **View Reports**.

   The **Overview** tab of the report is displayed, which shows a high-level summary of the file analysis.

6 Click the down-arrow on the **Report** tab and select a report to view.

   Key information about the analysis is displayed. The information in the report varies depending on the type of file that was analyzed.

7 (Optional) Click the icons on the **Report** tab to perform the following actions:

   - Export the report to other formats. Supported file formats are XML, JSON, and PDF
   - Download screenshots.
   - Download the packet capture (PCAP) file.

What to do next

To learn more about the information that is displayed on the **Overview** and **Report** tabs, see Analysis Report Details.

Analysis Report Details

Analysis report contains the detailed results of a file submission to the cloud. The **Overview** tab shows a high-level summary of the file analysis. The **Report** tab shows key information about the analysis that was performed on the file.

Overview Tab

The overview information is organized in the following sections.

Analysis Overview

This section provides a summary of the file analysis results. The following data is displayed:

- MD5 hash
- SHA1 hash
- SHA256 hash
- MIME type

- Submission timestamp

Threat Level

This section starts with a summary of the analysis findings.

For example: The file *md5_hash* was found to be Malicious.

After the summary, the following data is displayed:

Risk Assessment

- Maliciousness score: A score out of 100.
- Risk estimate: An estimate of the risk posed by the artifact.
  - High: The artifact represents a critical risk and must be addressed in priority. Such subjects are typically Trojan files or documents that contain exploits, leading to major compromises of the infected system. The risks are multiple: from information leakage to the system dysfunction. These risks are partially inferred from the type of activity detected. The score threshold for this category is usually ≥ 70.
  - Medium: The artifact represents a long-term risk and must be monitored closely. Such subjects can be a Web page containing suspicious content, potentially leading to drive-by attempts. They can also be adware or fake antivirus products that do not pose an immediate serious threat but can cause issues with the functioning of the system. The score threshold for this category is usually from 30–69.
  - Low: The artifact is benign and can be ignored. The score threshold for this risk estimate is usually below 30.
- Antivirus class: The antivirus or malware class to which the artifact belongs. For example, Trojan horse, worm, adware, ransomware, spyware, and so on.
- Antivirus family: The antivirus or malware family to which the artifact belongs. For example, valyria, darkside, and so on.

Analysis Overview

The data is sorted by severity, and includes the following fields:

- Severity: A score between 0–100 of the maliciousness of the activities detected during analysis of the artifact. Additional icons indicate in which operating systems the corresponding activity was observed during the analysis.
- Type: The types of activities detected during analysis of the artifact. These include:
  - Autostart: Ability to restart after a machine shutdown.
  - Disable: Ability to deactivate critical components of the system.
  - Evasion: Ability to evade analysis environment.
  - File: Suspicious activity on the file system.
  - Memory: Suspicious activity within the system memory.

- Network: Suspicious activity at the network level.
- Reputation: Known source or signed by a reputable organization.
- Settings: Ability to permanently alter critical system settings.
- Signature: Malicious subject identification.
- Steal: Ability to access and potentially leak sensitive information.
- Stealth: Ability to remain unnoticed by users or analysis systems.
- Silenced: Benign subject identification.

- Description: A description corresponding to each type of activity detected during analysis of the artifact.
- ATT&CK TACTICS: The MITRE ATT&CK stage or stages of an attack. Multiple tactics are separated by commas.
- ATT&CK TECHNIQUES: The observed actions or tools a malicious actor might use. Multiple techniques are separated by commas.

Additional Artifacts

This section lists additional artifacts (files and URLs) that were observed during the analysis of the submitted sample and that were in turn submitted for in-depth analysis. This section includes the following fields:

- Description: Describes the additional artifact.
- SHA1: The SHA1 hash of the additional artifact.
- Content Type: The MIME type of the additional artifact.
- Score: The maliciousness score of the additional artifact.

Decoded Command Line Arguments

If any PowerShell scripts were executed during the analysis, the system decodes these scripts, making its arguments available in a more human-readable form.

Third-party Tools

A link to a report on the artifact on VirusTotal portal.

Report Tab

Click the down-arrow on the **Report** tab and select a report to view. The information in the report varies depending on the type of file that was analyzed.

Analysis Information

This section contains the following key information about the analysis that the current report refers to:

- Analysis subject: The MD5 hash of the file.

- Analysis type: The type of analysis that was performed:
  - Dynamic analysis on Microsoft Windows 10: The analysis subject was run in a virtual Windows 10 environment using the VMware NSX® Network Detection and Response™ sandbox. The system monitors the file behavior and its interactions with the operating system looking for suspicious or malicious indicators.
  - Dynamic analysis on Microsoft Windows 7: The analysis subject was run in a virtual Windows 7 environment using the sandbox. The system monitors the file behavior and its interactions with the operating system looking for suspicious or malicious indicators.
  - Dynamic analysis in instrumented Chrome browser: The analysis subject (such as an HTML file or URL) was inspected using the instrumented browser, which is based on Google Chrome. The instrumented browser reproduces faithfully the behavior of the real browser and therefore is not easily fingerprinted by malicious content.
  - Dynamic analysis in emulated browser: The analysis subject (such as an HTML file or URL) was inspected using the emulated browser. The emulated browser can dynamically emulate different browser "personalities" (for example, changing its user-agent or varying the APIs that it exposes). This capability is useful when analyzing malicious content that targets specific browser types or versions. The drawback of this analysis type is that this browser is less realistic and can possibly be fingerprinted by malicious content.
  - Dynamic analysis in simulated file-viewer: The analysis subject (such as a PDF file) was inspected using the simulated file-viewer. The viewer can detect embedded contents and links.
  - Archive inflation: The analysis subject (an archive) was inflated, its contents were extracted and, if of appropriate type, was submitted for analysis.
- Password used: If available, the password that was used in the backend to successfully decrypt the sample, is provided.

View File Event Details in NSX Network Detection and Response UI

When files are extracted from the NSX Edges or the guest VM endpoints in the NSX, the generated file events are also sent to NSX Network Detection and Response application that is running in the cloud. In the NSX Network Detection and Response UI, you can correlate these file events with the other events in a Campaign, such as IDS events and Anomaly events.

The following procedure explains two methods to open the NSX Network Detection and Response UI in NSX Manager and view the file events.

Prerequisites

- NSX Malware Prevention and NSX Network Detection and Response features are activated on the NSX Application Platform.

- NSX Malware Prevention feature is activated on the ESXi host clusters or tier-1 gateways, or both, depending on your security requirements.

Procedure

1. From your browser, log in to an NSX Manager at https://*nsx-manager-ip-address*.
2. Click **Security**, and then in the left navigation pane, click **Malware Prevention**.

   The **Potential Malware** page is displayed.
3. Use any of the following methods to view the file event details in the NSX Network Detection and Response UI.
   - Method 1: On the **Potential Malware** page or the **All Files** page, expand a row to view the last inspection details of the file. Click the **Event Details** link. The **Event Profile** page opens in the NSX Network Detection and Response application.
   - Method 2: In the upper-right corner of the NSX Manager UI, click  , and then click **Network Detection and Response**. The **Dashboard** page of NSX Network Detection and Response is displayed. In the left navigation pane, click **Events** and search the events that you want to view. Click the event to see more details in the **Event Summary** pane.

View Campaign Details in NSX Network Detection and Response UI

From the **Malware Prevention** dashboard, you can switch to the **Campaigns** page in the NSX Network Detection and Response UI and monitor the campaigns that NSX Network Detection and Response has detected in your network.

The following procedure explains multiple methods to open the NSX Network Detection and Response UI in NSX Manager and view the campaign details.

Prerequisites

- NSX Malware Prevention and NSX Network Detection and Response features are activated on the NSX Application Platform.
- NSX Malware Prevention feature is activated on the ESXi host clusters or tier-1 gateways, or both, depending on your security requirements.

Procedure

1. From your browser, log in to an NSX Manager at https://*nsx-manager-ip-address*.
2. Click **Security**, and then in the left navigation pane, click **Malware Prevention**.

   The **Potential Malware** page is displayed.

3 Use any of the following methods to view campaigns in the NSX Network Detection and Response UI.

- Method 1: On the **Potential Malware** page or the **All Files** page, expand a row to view the last inspection details of the file. Click the **Campaigns** link. This link appears only when campaigns are available for the detected malware.
- Method 2: In the upper-right corner of the NSX Manager UI, click  , and then click **Network Detection and Response**. The **Dashboard** page of NSX Network Detection and Response is displayed.

  The **Active campaigns in my network** widget on the **Dashboard** page provides an overview of the campaigns that NSX Network Detection and Response has identified and that are currently active in your network. This widget helps you to focus your attention on the most critical campaigns for immediate action. It displays statistics for All Active Campaigns, Open High Impact Campaigns, In-Progress High Impact Campaigns, and Hosts Affected.

  To see more details about these campaigns, click **Go to campaigns overview** at the bottom-left corner of the widget.
- Method 3: In the NSX Manager UI, navigate to **Security > Security Overview > Threat Detection & Response > Campaigns**.

For more information about campaigns, see the *Using and Managing VMware NSX Intelligence* documentation at https://docs.vmware.com/en/VMware-NSX-Intelligence/index.html.

Monitoring IDS/IPS Events

You can monitor events and view data of the last 14 days.

To view intrusion events, navigate to **Security > IDS/IPS** . You can filter the events based on the following criteria:

- Filter criteria. Select from the following options:

Filter Criteria	Description
Attack Target	Target of the attack.
Attack Type	Type of attack, such as trojan horse, or denial of service (DoS).
CVSS	Common Vulnerability Score (filter based on a score above a set threshold).
Gateway Name	The gateway name on which the event was registered.
IP Address	IP address on which the event was registered.
Product Affected	Vulnerable product or (version), such as Windows XP or Web_Browsers.

Filter Criteria	Description
Signature ID	Unique ID of the signature rule.
VM Name	The VM (based on logical port) on which the event was registered.

- Traffic: Select from the following options:
  - All traffic
  - Distributed only
  - Gateway only
- Signature actions: Select from the following options:
  - Show all signatures
  - Dropped (Prevented)
  - Rejected (Prevented)
  - Alert (Detect Only)
- Severity rating: Select from the following options:
  - Critical
  - High
  - Medium
  - Low
  - Suspicious

You can toggle the **Timeline** button to view or hide the timeline graph that is based on severity ratings. The graph presents events that occurred over a selected time span. You can zoom in to the specific time window on this graph to view details of signatures of the related events that happened during the time window.

On the timeline graph, colored dots indicate the unique type of intrusion events and can be clicked for details. The size of the dot indicates the number of times an intrusion event has been seen. A blinking dot indicates that an attack is ongoing. Point to a dot to see the attack name, number of attempts, first occurrence, and other details.

- Red dots - represent critical severity signature events.
- Orange dots - represent high severity signature events.
- Yellow dots - represent medium severity signature events.
- Gray dots - represent low severity signature events.
- Purple - represent suspicious severity signature events.

All the intrusion attempts for a particular signature are grouped and plotted at their first occurrence.

Click the arrow next to an event to view details.

Detail	Description
Impact Score	Impact score is a combined value of risk score (the severity of the threat) and the confidence score (strength of the detection being correct).
Severity	Signature severity of the intrusion.
Last Detected	This is the last time the signature was fired.
Details	Brief description of what the signature is targeting.
Users Affected	Number of users who were impacted by the event.
Workloads	Number of workloads affected. Click to view affected workload details.
CVE Details	CVE reference of the vulnerability targeted by the exploit.
CVSS	Common Vulnerability Score of the vulnerability targeted by the exploit.
Intrusion Event Details (latest occurrence) - Source	IP address of the attacker and source port used.
Intrusion Event Details (latest occurrence) - Gateway	Edge node details that contain the workload on which the event was registered.
Intrusion Event Details (latest occurrence) - Hypervisor	Transport node details that contain the workload on which the event was registered.
Intrusion Event Details (latest occurrence) - Target	IP address of the victim and destination port used.
Attack Direction	Client-Server or Server-Client.
Attack Target	Target of the attack.
Attack Type	Type of attack, such as trojan horse, or denial of service (DoS).
Product Affected	Illustrates what product is vulnerable to the exploit.
Total Events	Total number of intrusion attempts for the event.
Intrusion Activity	Displays the total number of times this particular IDS signature was triggered, the most recent occurrence, and the first occurrence.
Service	Protocol information associated with the event.
Signature ID	Unique ID of the IDS signature.
Signature Revision	The revision number of the IDS signature.
Mitre Technique	MITRE ATT&CK technique describing the detected activity.
Mitre Tactic	MITRE ATT&CK tactic describing the detected activity.
Associated IDS Rule	Clickable link to the configured IDS Rule which resulted in this event.

To view full intrusion history, click the **View Full Event History** link. A window opens with the following details:

Detail	Description
Time Detected	This is the last time the signature was fired.
Traffic Type	This could be Distributed or Gateway. Distributed indicates East-West traffic flow and Gateway indicates North-South traffic flow.
Workloads/IPs Affected	Number of virtual machines or IP addresses which has hit the given attack or vulnerability for a given traffic flow.
Attempts	Number of intrusion attempts made for an attack or vulnerability during a given traffic flow.
Source	IP address of the attacker.
Destination	IP address of the victim.
Protocol	Traffic protocol of the detected intrusion.
Rule	Rule to which the signature belongs (through the profile).
Profile	Profile to which the signature belongs.
Action	Any of the following actions that was triggered against the event: ■ Drop ■ Reject ■ Alert

You can also filter intrusion history based on the following criteria:

- Action
- Destination IP
- Destination Port
- Protocol
- Rule
- Source IP
- Source Port
- Traffic Type

Remote logging

NSX components write to log files in the directory `/var/log`. On NSX appliances, NSX syslog messages conform with RFC 5424. On ESXi hosts, syslog messages conform with RFC 3164.

There are three event log files in the `/var/log/nsx-idps` folder on ESXi hosts:

- fast log - Contains internal logging of nsx-idps process events, with limited information and is used only for debugging purposes.
- nsx-idps-log - Contains general nsx-idps process logs with basic information and errors about the process workflow.

- nsx-idps-events.log - Contains detailed information about events (all alerts/drops/rejects) with NSX metadata.

To enable the sending of NSX IDS/IPS logs to a central log repository, use the following APIs.

Run the following API to query the current settings.

```
GET https://<Manager-IP>/api/v1/global-configs/IdsGlobalConfig
```

```
{
   ...
   "global_idsevents_to_syslog_enabled": false,
   ...
 }
```

If the *global_idsevents_to_syslog.enabled* variable is set to false, run the following API to set it to true.

```
PUT https://<Manager-IP>/api/v1/global-configs/IdsGlobalConfig
```

```
Result:
    {
      "global_idsevents_to_syslog_enabled": true,
      "_revision": <revision number from GET>
     }
```

These events are exported directly from ESXi hosts so ensure remote syslog is configured on the ESXi host.

You must also ensure that the NSX manager and ESXi hosts are also setup to forward syslog messages to the central log repository.

For more information about configuring remote logging, see Configure Remote Logging and all related information under the section Log Messages and Error Codes.

Administering NSX Malware Prevention

You can upgrade or delete the NSX Malware Prevention feature by using the **NSX Application Platform** page in the **Security** tab.

Upgrading NSX Malware Prevention

Upgrading NSX Malware Prevention currently occurs only when NSX Application Platform is upgraded.

See the "Upgrade NSX Application Platform" topic in the *Deploying and Managing the VMware NSX Application Platform* documentation.

Delete NSX Malware Prevention

If for some reason you want to delete NSX Malware Prevention, use the steps described in this section.

Caution When you delete the NSX Malware Prevention feature, the system deletes all the data analytics that have been gathered, along with the feature. The action is permanent and results in loss of previously collected data.

Prerequisites

- NSX Application Platform must be in a good state and there are no active alarms.
- You must have NSX **Enterprise Administrator** privileges.
- A valid NSX Data Center edition license is in effect for your NSX Manager session. See System Requirements for NSX IDS/IPS and NSX Malware Prevention for license information.

Procedure

1. From your browser, log in with **Enterprise Administrator** privileges to an NSX Manager at `https://nsx-manager-ip-address`.
2. In the NSX Manager UI, select **System > NSX Application Platform** in the Configuration section.
3. Navigate to the Features section and in the NSX Malware Prevention feature card, click **Actions** and select **Delete**.
4. Click **Delete** in the **Delete Malware Prevention** dialog box.

Results

After a successful deletion, the NSX Malware Prevention feature card returns to a ready-to-activate state.

Troubleshooting NSX Malware Prevention

Use the information in this chapter to understand log messages, resolve syslog issues, and troubleshoot common problems that can occur with the NSX Malware Prevention feature.

Collecting Logs

NSX Malware Prevention feature runs on NSX Edges, service virtual machine (on ESXi hosts), and NSX Application Platform. The product logs generated on NSX Edges and service virtual machines conform to the RFC 5424 log message standard. NSX Malware Prevention is supported only on ESXi hosts.

Log Messages

On NSX appliances, syslog messages conform to the RFC 5424 standard. Additional product logs are written to the `/var/log` directory.

Remote logging is also supported. To consume NSX Malware Prevention feature logs, you can configure NSX Edges and NSX Application Platform to send logs to a remote log server.

Configure Remote Logging on NSX Edge

You must configure remote logging on each NSX Edge node individually. To configure the remote logging server on an NSX Edge node by using the NSX CLI, see Configure Remote Logging.

To configure the remote logging server on an NSX Edge node by using the NSX Manager UI, see Add Syslog Servers for NSX Nodes.

Configure Remote Logging on NSX Application Platform

To send NSX Application Platform log files to an external log server, you must run a REST API.

Use the following NSX Application Platform Cluster API to add an external syslog server to stream the logs to:

```
POST https://{api_host}/napp/api/v1/platform/support-bundle/external-syslog-server
```

For information about the parameters of this REST API along with sample request body, response, and code samples, see the VMware Developer Documentation portal.

Configure Remote Logging on NSX Malware Prevention Service Virtual Machine

This functionality is currently not supported. However, as a workaround, you can collect log files from each NSX Malware Prevention service virtual machine (SVM) by logging on to the SVM with an SSH connection.

SSH access to the **admin** user of the SVM is key-based (public-private key pair). A public key is needed when you are deploying the service on an ESXi host cluster, and a private key is needed when you want to start an SSH session to the SVM.

For more information see, Log in to the NSX Malware Prevention Service Virtual Machine

After logging in to the SVM, use the `sftp` or the `scp` command to collect the logs from the `/var/log` directory at that particular time.

More Information about Logging

See Log Messages and Error Codes.

Interpret NSX Malware Prevention Event Logs

The following event log is generated by the `sa-events-processor` microservice, which is a pod that runs on the NSX Application Platform.

Example:

```
{"log":"{\"log\":\"\\u001b[37m2022-06-01T01:42:58,725\\u001b[m \\u001b[32mINFO \\u001b[m[\
\u001b[1;34mfileEventConsumer-1\\u001b[m] \
\u001b[1;33mc.v.n.s.e.k.EventsProcessorConsumerService\\u001b[m: SECURITY [nsx@6876 comp=\\
\"nsx-manager\\\" level=\\\"INFO\\\" subcomp=\\\"manager\\\"] Event number 2 received from
topic: ams-file-seen-events partition: 0 and offset: 212640 is: FileInspectionEvent(id=0,
```

```
sha256=29fbd4604acb1da497e8127cd688bf2614f565fc4d4c808989df41c4a6fb924d,
sha1=549cb3f1c85c4ef7fb06dcd33d68cba073b260ec, md5=65b9b68668bb6860e3144866bf5dab85,
fileName=drupdate.dll, fileType=PeExeFile, fileSize=287024, inspectionTime=1654047770305,
clientPort=0, clientIp=null, clientFqdn=null,
clientVmId=500cd1b6-96b6-4567-82f4-231a63dead81, serverPort=0, serverIp=null,
serverFqdn=null, serverVmId=null, applicationProtocol=null, submittedBy=SYSTEM,
isFoundByAsds=true, isBlocked=false, allowListed=false, verdict=BENIGN, score=0,
analystUuid=null, submissionUuid=null, tnId=38c58796-9983-4a41-b9f2-dc309bd3458d,
malwareClass=null, malwareFamily=null, errorCode=null, errorMessage=null, nodeType=1,
gatewayId=, analysisStatus=COMPLETED, followupEvent=false, httpDomain=null, httpMethod=null,
path=null, referer=null, userAgent=null, contentDispositionFileName=null, isFileUpload=false,
startTime=1654047768828, endTime=1654047768844, ttl=1654220570304)\
\n\",\"stream\":\"stdout\",\"time\":\"2022-06-01T01:42:58.725811209Z\"}","log_processed":
{"log":"\u001b[37m2022-06-01T01:42:58,725\u001b[m \u001b[32mINFO
\u001b[m[\u001b[1;34mfileEventConsumer-1\u001b[m]
\u001b[1;33mc.v.n.s.e.k.EventsProcessorConsumerService\u001b[m: SECURITY [nsx@6876 comp=\"nsx-
manager\" level=\"INFO\" subcomp=\"manager\"] Event number 2 received from topic: ams-file-
seen-events partition: 0 and offset: 212640 is: FileInspectionEvent(id=0,
sha256=29fbd4604acb1da497e8127cd688bf2614f565fc4d4c808989df41c4a6fb924d,
sha1=549cb3f1c85c4ef7fb06dcd33d68cba073b260ec, md5=65b9b68668bb6860e3144866bf5dab85,
fileName=drupdate.dll, fileType=PeExeFile, fileSize=287024, inspectionTime=1654047770305,
clientPort=0, clientIp=null, clientFqdn=null,
clientVmId=500cd1b6-96b6-4567-82f4-231a63dead81, serverPort=0, serverIp=null,
serverFqdn=null, serverVmId=null, applicationProtocol=null, submittedBy=SYSTEM,
isFoundByAsds=true, isBlocked=false, allowListed=false, verdict=BENIGN, score=0,
analystUuid=null, submissionUuid=null, tnId=38c58796-9983-4a41-b9f2-dc309bd3458d,
malwareClass=null, malwareFamily=null, errorCode=null, errorMessage=null, nodeType=1,
gatewayId=, analysisStatus=COMPLETED, followupEvent=false, httpDomain=null, httpMethod=null,
path=null, referer=null, userAgent=null, contentDispositionFileName=null, isFileUpload=false,
startTime=1654047768828, endTime=1654047768844,
ttl=1654220570304)","stream":"stdout","time":"2022-06-01T01:42:58.725811209Z"},"kubernetes":
{"pod_name":"sa-events-processor-55bcfcc46d-4jftf","namespace_name":"nsxi-
platform","pod_id":"305953f7-836b-4bbb-
ba9e-00fdf68de4ae","host":"worker03","container_name":"sa-events-
processor","docker_id":"93f81f278898e6ce3e14d9a37e0e10a502c46fe53c9ad61680aed48b94f7f8bf","con
tainer_hash":"projects.registry.vmware.com/nsx_application_platform/clustering/sa-events-
processor@sha256:b617f4bb9f3ea5767839e39490a78169f7f3d54826b89638e4a950e391405ae4","container_
image":"projects.registry.vmware.com/nsx_application_platform/clustering/sa-events-
processor:19067767"}}
```


Note This example event log is only for illustrative purposes. The format and content might change across major NSX versions.

In this sample event log, observe that apart from the standard log attributes, such as `date` (2022-06-01T00:42:58,326), `log level` (INFO), and filterable attributes, such as `module` (SECURITY), `container_name` (sa-events-processor), additional attributes are present in a JSON style format. The following table lists these additional attributes.

Key	Sample Value	Key	Sample Value
id	0	score	0
sha256	29fbd4604acb1da497e8127cd688bf2614f565fc4d4c808989df41c4a6fb924d	analystUuid	null
sha1	549cb3f1c85c4ef7fb06dcd33d68cba073b260ec	submissionUuid	null
md5	65b9b68668bb6860e3144866bf5dab85	tnId	38c58796-9983-4a41-b9f2-dc309bd3458d
fileName	drupdate.dll	malwareClass	null
fileType	PeExeFile	malwareFamily	null
fileSize	287024	errorCode	null
inspectionTime	1654047770305	errorMessage	null
clientPort	0	nodeType	1
clientIP	null	gatewayId	
clientFqdn	null	analysisStatus	COMPLETED
clientVmId	500cd1b6-96b6-4567-82f4-231a63dead81	followupEvent	false
serverPort	0	httpDomain	null
serverIp	null	httpMethod	null
serverFqdn	null	path	null
serverVmId	null	referer	null
applicationProtocol	null	userAgent	null
submittedBy	SYSTEM	contentDispositionFileName	null
isFoundByAsds	true	isFileUploaded	false
isBlocked	false	startTime	1654047768828
allowListed	false	endTime	1654047768844
verdict	BENIGN	ttl	1654220570304

Troubleshoot Syslog Issues

If the remote log server that you configured is unable to receive logs, see Troubleshooting Syslog Issues.

Collect Support Bundles

- To collect support bundles for Management Nodes, NSX Edges, and Hosts, see Collect Support Bundles.
- To collect support bundles for NSX Application Platform, see the *Deploying and Managing the VMware NSX Application Platform* documentation at https://docs.vmware.com/en/VMware-NSX-T-Data-Center/index.html.

Troubleshooting NSX Malware Prevention Service Virtual Machine Problems

Use the information in this topic to debug issues that are associated with NSX Distributed Malware Prevention service deployment, health status of service instances, ESXi Agencies, and other issues.

Verify ESX Agent Manager Health Status

To verify whether the health status of vSphere ESX Agent Manager (EAM) is normal, do these steps:

1 In the vSphere Client, navigate to **Administration > vCenter Server Extensions**. Click **vSphere ESX Agent Manager**.

2 Click the **Configure** tab.

This page shows the health status of ESX Agencies on the hosts for the NSX Malware Prevention solution and issues (if any) that are detected for the Agencies.

ESXi Agency Manager (EAM) service must be up and running. Verify whether the following URL is accessible.

`https://vCenter_Server_IP_Address/eam/mob`

Replace *vCenter_Server_IP_Address* with the IP address of the VMware vCenter in your network.

Verify Connectivity of Port Groups, Interfaces, and Context Multiplexer

Do the following steps in the vSphere Client:

- Select the name of the service virtual machine, and then click the **Networks** tab. Verify that the **vmservice-vhsield-pg** Port Group is listed.
- Right-click the service virtual machine name, and click **Edit Settings**. On the **Virtual Hardware** page, verify that network adapter 1 and network adapter 2 are connected. Network adapter 1 connects the SVM to the Management network, and network adapter 2 connects the SVM to the **vmservice-vshield-pg** Port Group, which NSX has autocreated during the service deployment. Network adapter 2 is the control interface of the SVM that is used for communication between the Context Multiplexer (MUX) and the SVM. For NSX Malware Prevention SVM, the control interface IP is 169.254.1.22.

Context Multiplexer service must be running on each ESXi host. To verify whether the `nsx-context-mux` service is running on the host, log in to the CLI of each ESXi host as a root user and run the following CLI command:

```
# /etc/init.d/nsx-context-mux  status
```

If the service is not running, start or restart the service with the following CLI command:

```
/etc/init.d/nsx-context-mux start
```

Or

```
/etc/init.d/nsx-context-mux restart
```

Note It is safe to restart this service during production hours because restarting the service does not have a significant impact. The service restarts in a couple of seconds.

Resolve ESX Agent Manager Issues

The ESX Agent Manager notifies NSX Manager about error details when it detects issues in the ESX Agencies. You can click **Resolve** in the NSX Manager UI to resolve the issues. The following table describes the ESX Agent Manager issues.

Issue	Category	Description	Resolution
Cannot Access Agent OVF	VM Not Deployed	An agent virtual machine is expected to be deployed on a host, but the agent virtual machine cannot be deployed because the ESXi Agent Manager is unable to access the OVF package for the agent. It might happen because the web server providing the OVF package is down. The web server is often internal to the solution that created the Agency.	ESXi Agency Manager (EAM) service retries the OVF download operation. Click **Resolve**.
Incompatible Host Version	VM Not Deployed	An agent virtual machine is expected to be deployed on a host. However, because of compatibility issues the agent was not deployed on the host.	Upgrade either the host or the solution to make the agent compatible with the host. Check the compatibility of the SVM. Click **Resolve**.
Insufficient Resources	VM Not Deployed	An agent virtual machine is expected to be deployed on a host. However, ESXi Agency Manager (EAM) service did not deploy the agent virtual machine because the host has less CPU or memory resources.	ESXi Agency Manager (EAM) service attempts to redeploy the virtual machine. Ensure that CPU and memory resources are available. Check the host and free up some resources. Click **Resolve**.
Insufficient Space	VM Not Deployed	An agent virtual machine is expected to be deployed on a host. However, the agent virtual machine was not deployed because the agent datastore on the host did not have enough free space.	ESXi Agency Manager (EAM) service attempts to redeploy the virtual machine. Free up some space on the datastore. Click **Resolve**.

Issue	Category	Description	Resolution
No Agent VM Network	VM Not Deployed	An agent virtual machine is expected to be deployed on a host, but the agent cannot be deployed because the agent network has not been configured on the host.	Add one of the networks listed in custom agent VM network to the host. The issue resolves automatically after the datastore is available.
OVF Invalid Format	VM Not Deployed	An Agent virtual machine is expected to be provisioned on a host, but it failed to do so because the provisioning of the OVF package failed. The provisioning is unlikely to succeed until the solution that provides the OVF package has been upgraded or patched to provide a valid OVF package for the agent virtual machine.	ESXi Agency Manager (EAM) service attempts to redeploy the SVM. Ensure that a valid OVF package is used for service deployment. Click **Resolve**.
Missing Agent IP Pool	VM Powered Off	An agent virtual machine is expected to be powered on, but the agent virtual machine is powered off because there are no IP addresses defined on the agent's virtual machine network.	Define the IP address on the virtual machine network. Click **Resolve**.
No Agent VM Datastore	VM Powered Off	An agent virtual machine is expected to be deployed on a host, but the agent cannot be deployed because the agent datastore has not been configured on the host.	Add one of the datastores listed in custom agent VM datastore to the host. The issue resolves automatically after the datastore is available.
No Custom Agent VM Network	No Agent VM Network	An agent virtual machine is expected to be deployed on a host, but the agent cannot be deployed because the agent network has not been configured on the host.	Add the host to one of the networks listed in a custom agent VM network. The issue resolves automatically after a custom VM network is available.
No Custom Agent VM Datastore	No Agent VM Datastore	An agent virtual machine is expected to be deployed on a host, but the agent cannot be deployed because the agent datastore has not been configured on the host.	Add the host to one of the datastores listed in a custom agent VM datastore. The issue resolves automatically.
Orphaned DvFilter Switch	Host Issue	A dvFilter switch exists on a host but no agents on the host depend on dvFilter. It happens if a host is disconnected when an agency configuration changed.	Click **Resolve**. ESXi Agency Manager (EAM) service attempts to connect the host before the agency configuration is updated.

Issue	Category	Description	Resolution
Unknown Agent VM	Host Issue	An agent virtual machine has been found in the vCenter Server inventory that does not belong to any agency in this vSphere ESX Agent Manager server instance.	Click **Resolve**. ESXi Agency Manager (EAM) service attempts to place the virtual machine to the inventory it belongs to.
OVF Invalid Property	VM Issue	An agent virtual machine must be powered on, but an OVF property is either missing or has an invalid value.	Click **Resolve**. ESXi Agency Manager (EAM) service attempts to reconfigure the correct OVF property.
VM Corrupted	VM Issue	An agent virtual machine is corrupt.	Click **Resolve**. ESXi Agency Manager (EAM) service attempts to repair the virtual machine.
VM Orphaned	VM Issue	An agent virtual machine exists on a host, but the host is no longer part of scope for the agency. It happens if a host is disconnected when the agency configuration is changed.	Click **Resolve**. ESXi Agency Manager (EAM) service attempts to connect the host back to the agency configuration.
VM Deployed	VM Issue	An agent virtual machine is expected to be removed from a host, but the agent virtual machine has not been removed. The specific reason why vSphere ESX Agent Manager was unable to remove the agent virtual machine, such as the host is in maintenance mode, powered off or in standby mode.	Click **Resolve**. ESXi Agency Manager (EAM) service attempts to remove the agent virtual machine from the host.
VM Powered Off	VM Issue	An agent virtual machine is expected to be powered on, but the agent virtual machine is powered off.	Click **Resolve**. ESXi Agency Manager (EAM) service attempts to power on the virtual machine.
VM Powered On	VM Issue	An agent virtual machine is expected to be powered off, but the agent virtual machine is powered on.	Click **Resolve**. ESXi Agency Manager (EAM) service attempts to power off the virtual machine.
VM Suspended	VM Issue	An agent virtual machine is expected to be powered on, but the agent virtual machine is suspended.	Click **Resolve**. ESXi Agency Manager (EAM) service attempts to power on the virtual machine.
VM Wrong Folder	VM Issue	An agent virtual machine is expected to be in a designated agent virtual machine folder, but is found in a different folder.	Click **Resolve**. ESXi Agency Manager (EAM) service attempts to place the agent virtual machine to the designated folder.

Issue	Category	Description	Resolution
VM Wrong Resource Pool	VM Issue	An agent virtual machine is expected to be located in a designated agent virtual machine resource pool, but is found in a different resource pool.	Click **Resolve**. ESXi Agency Manager (EAM) service attempts to place the agent virtual machine to a designated resource pool.
VM Not Deployed	Agent Issue	An agent virtual machine is expected to be deployed on a host, but the agent virtual machine has not been deployed. Specific reasons why ESXi Agent Manager was unable to deploy the agent, such as being unable to access the OVF package for the agent or a missing host configuration. This issue can also happen if the agent virtual machine is explicitly deleted from the host.	Click **Resolve** to deploy the agent virtual machine.

Resolve NSX Manager Issue

Issue

Unable to allocate static IP addresses from the IP Pool.

Description

Either the IP addresses from the pool are exhausted or there are no more IP addresses left to allocate.

Resolution

Fix the IP Pool problem, and then click **Resolve** to fix the issue.

Verify Health Status of Service Instances

NSX Manager receives the health status details of each service instance. The latest timestamp when the health status is received is shown in the NSX Manager UI. You might have to refresh the **Service Instances** page a few times to retrieve the latest health status.

Health of a service instance on an ESXi host depends on the following factors:

Solution status

Status of the NSX Distributed Malware Prevention solution that is running on an SVM. Up status indicates that the solution is correctly running.

Connectivity between NSX Guest Introspection agent and Context engine

Status is Up when NSX Guest Introspection agent (Context Multiplexer) is connected to the NSX Ops agent, which includes the Context engine. The Context Multiplexer forwards health information of SVMs to the Context engine. The MUX and the NSX Ops agent also share SVM-VM configuration between each other to know which workload VMs are protected by the SVM.

Service VM protocol version

Transport protocol version used internally for troubleshooting issues.

NSX Guest Introspection agent information

Represents protocol version compatibility between NSX Guest Introspection agent and SVM.

To view the health status of service instances, do these steps in NSX Manager:

1 Navigate to **System > Service Deployments > Service Instances**.

2 In the **Health Status** column, click the icon next to Up or Down.

View Alarms in NSX Manager

Alarms are displayed on the **Alarms** page of the NSX Manager UI for the following situations:

- Connectivity between NSX Context Multiplexer and NSX Malware Prevention SVM is down.
- NSX Context Multiplexer is down or reboots.

You can also view alarms on the **Service Instances** page at **System > Service Deployments > Service Instances**.

To view alarms about the health of the NSX Malware Prevention feature, do these steps:

1 In NSX Manager, navigate to the **Alarm Definitions** page by clicking **Home > Alarms > Alarm Definitions**.

2 Click in the **Filter by Name, Path, and more** text area, and then click **Feature**.

3 Select the **Malware Prevention Health** check box.

For documentation about the NSX Malware Prevention health events, see the NSX Event Catalog.

View Component Issues on Security Overview Dashboard

The Malware Prevention widget on the **Security Overview** dashboard shows issues when any of the components in the NSX Distributed Malware Prevention service is down or not working. To view this UI widget in the NSX Manager UI, navigate to **Security > Security Overview > Configuration**.

For example:

- The Bar chart shows an issue when the Security Hub on the NSX Malware Prevention service virtual machine (SVM) is down. Point to the bar to view the following details:
  - Number of NSX Malware Prevention SVMs that are impacted.

- Number of workload VMs on the host that have lost malware security protection due to the Security Hub going down.

- The Donut chart shows the following details:
  - Number of workload VMs where the NSX File Introspection driver is running.
  - Number of workload VMs where the NSX File Introspection driver is not running.

  For both these metrics, only the workload VMs on the host clusters that are activated for NSX Distributed Malware Prevention are considered.

Name the Key Pairs Correctly for Easy Identification

SSH access to the **admin** user of the SVM is key-based (public-private key pair). A public key is needed when you are deploying the service on an ESXi host cluster, and a private key is needed when you want to start an SSH session to the SVM.

NSX Distributed Malware Prevention service deployment is done at the level of a host cluster. So, a key pair is tied to a host cluster. You can create either a new public-private key pair for a service deployment on each cluster, or use a single key pair for service deployments on all the clusters.

If you plan to use a different public-private key pair for service deployment on each cluster, ensure that the key pairs are named correctly for easy identification.

A good practice is to identify each service deployment with a "compute cluster id" and specify the cluster id in the name of the key pair. For example, let us assume that the cluster id is "1234-abcd". For this cluster, you can specify the service deployment name as "MPS-1234-abcd", and name the key pair to access this service deployment as "id_rsa_1234_abcd.pem". This practice makes it easy for you to maintain and associate keys for each service deployment.

Important Store the private key securely. Loss of the private key can lead to a loss of SSH access to the NSX Malware Prevention SVM.

If the private key is lost, the SVM continues to function without any issues, but you cannot log in to the SVM and download the log file for troubleshooting purposes.

Collect Support Bundle and NSX Malware Prevention SVM Log

For a detailed troubleshooting of the following components that contribute to NSX Distributed Malware Prevention service, you can collect support bundles for analysis or send them to VMware Support:

- NSX Manager appliances
- ESXi hosts
- VMware Tools on workload VMs
- NSX Malware Prevention SVM

To collect a support bundle that contains log files for ESXi hosts and NSX Manager appliances, use the support bundle feature in NSX. For instructions about collecting a support bundle in NSX, see Collect Support Bundles.

To collect a support bundle that contains log files for components and services running on VMware vCenter, see the *vCenter Server Configuration* documentation. For example, you can collect log files for VMware Tools with the VMware vCenter support bundle.

To collect log files for the NSX Malware Prevention SVM, start a remote SSH session to the SVM with the private key of the SVM. For more information, see Log in to the NSX Malware Prevention Service Virtual Machine.

The log files are available on the SVM at `/var/log`. If multiple syslog files are available at this location, they are compressed and stored at the same path.

NSX Network Detection and Response

To defend your network environment against MITRE ATT&CK techniques, VMware NSX® Network Detection and Response™ uses signals received from the network traffic analysis, intrusion detection and prevention, and network sandboxing engines available in NSX Data Center. NSX Network Detection and Response provides a cloud-based architecture that enables your security operations team to gain comprehensive visibility into the traffic that crosses the perimeter (north/south), and traffic that moves laterally inside the perimeter (east/west) of your network.

This security feature is part of the VMware NSX® Advanced Threat Prevention solution.

Getting Started with NSX Network Detection and Response

This section provides an overview of the NSX Network Detection and Response feature. To use the NSX Network Detection and Response functionalities, understand the system requirements and the other NSX features that you must configure.

Overview of NSX Network Detection and Response

The main objective of the NSX Network Detection and Response feature is to collect key abnormal activity or malicious events from every activated event source in your NSX environment.

Collected Events

NSX Network Detection and Response submits any collected events that require further analysis to the VMware NSX® Advanced Threat Prevention cloud service for correlation and visualization. You can view and manage the analysis results using the NSX Network Detection and Response user interface (UI).

NSX Network Detection and Response correlates events that it determines to be related into campaigns. It organizes threat events in a campaign into a timeline that is available for a security analyst to view and triage using the NSX Network Detection and Response UI.

Event Types and Event Sources

The following table lists the event types that NSX Network Detection and Response can collect and the sources that generate those events. In order for any of the event source to send the events to NSX Network Detection and Response, you must activate the corresponding NSX feature mentioned for the event type.

Event Type	Events Source
Malicious file events	Edge appliance, if you activate the VMware NSX® Malware Prevention feature.
IDS events	Distributed IDS, if you activate the Distributed NSX IDS/IPS feature.
Network traffic anomaly events	VMware NSX® Intelligence™, if activated, and if you turn on the NSX Suspicious Traffic detectors.

Important To maximize the NSX Network Detection and Response feature, activate one or more of the NSX features whose events it consumes. Although you can activate the NSX Network Detection and Response feature on its own, if you do not activate any of the NSX features mentioned in the previous table, NSX Network Detection and Response does not have any events to analyze and, thus, cannot give any of the benefits it has to offer.

Activating and Using the Feature

Before you can start using the NSX Network Detection and Response feature, you must meet specific license requirements and software requirements, and you must activate the feature. To start using NSX Network Detection and Response to manage the different event types that you can monitor in your NSX environment, you must also activate and configure the corresponding NSX features.

For more information on the next steps, see NSX Network Detection and Response Activation and Usage Workflow.

Activating Other NSX Features

For information about how to activate and configure the NSX features whose detection events NSX Network Detection and Response consumes, refer to the following table.

NSX Feature to Activate	Documentation Name and Location	Topic Title
NSX IDS/IPS	*NSX Administration Guide* for version 3.2 or later.	Getting Started with NSX IDS/IPS and NSX Malware Prevention
NSX Malware Prevention	*NSX Administration Guide* for version 3.2 or later.	Activate NSX Malware Prevention

NSX Feature to Activate	Documentation Name and Location	Topic Title
NSX Intelligence	*Activating and Upgrading VMware NSX Intelligence* for version 3.2 or later delivered with the VMware NSX Intelligence Documentation set.	Activate NSX Intelligence
NSX Suspicious Traffic	*Using and Managing VMware NSX Intelligence* for version 3.2 or later delivered with the VMware NSX Intelligence Documentation set.	Activate the NSX Suspicious Traffic Detectors

NSX Network Detection and Response Activation and Usage Workflow

To track your progress in activating NSX Network Detection and Response and to guide you on how to begin using the feature, use the following checklist.

Perform steps 1-5 in the order they are listed. Perform the remaining steps depending on your needs.

1 Install NSX 3.2 or later.

   See the installation workflow details in the *NSX Installation Guide* documentation delivered with the VMware NSX Documentation set.

2 Ensure you have reviewed and met the NSX Network Detection and Response system requirements listed in System Requirements for NSX Network Detection and Response.

3 Deploy the NSX Application Platform using the NSX Manager 3.2 or later user interface.

   NSX Network Detection and Response is an application hosted on the NSX Application Platform. See the *Deploying and Managing the VMware NSX Application Platform* documentation delivered with the VMware NSX Documentation set.

4 Activate the NSX Network Detection and Response feature. See Activate NSX Network Detection and Response.

5 To learn more about working with the NSX Network Detection and Response functionalities, see Working with the NSX Network Detection and Response Application.

6 Start using NSX Network Detection and Response and NSX Suspicious Traffic to view details about anomalous or suspicious network traffic events. This requires a separate NSX Intelligence activation and configuration, and also turning on NSX Suspicious Traffic detectors.

   For details on how to detect suspicious network activities using NSX Intelligence 3.2 or later, see the topics in the "Detecting Suspicious Network Traffic in NSX" section of the *Using and Managing VMware NSX Intelligence* documentation delivered with the VMware NSX Intelligence Documentation set.

7 Start viewing details about malicious file events generated by the NSX Malware Prevention feature. This requires a separate activation and configuration of the NSX Malware Prevention feature.

For details on how to activate and configure the NSX Malware Prevention features, see Activate NSX Malware Prevention .

8 Work with IDS event details using the NSX Distributed IDS/IPS feature. This requires separate NSX Distributed IDS/IPS feature activation and configuration.

See the details on how to activate and configure NSX Distribtued IDS/IPS in Getting Started with NSX IDS/IPS and NSX Malware Prevention.

System Requirements for NSX Network Detection and Response

To use the NSX Network Detection and Response feature, you must prepare your NSX environment so that it meets the specific license and software requirements.

License Requirements

You must have one of the following license in effect during your NSX Manager session. The following lists the various NSX licenses that support the NSX Network Detection and Response feature.

Base SKU License	Add-on SKU License
NSX-T Evaluation	None required
NSX Data Center Evaluation	None required
NSX Advanced Threat Prevention (Only applicable for customers who have previously purchased the license.)	None required
One of the following: ■ NSX Distributed Firewall with Threat Prevention ■ NSX Distributed Firewall ■ NSX Advanced ■ NSX Enterprise Plus	NSX Advanced Threat Prevention for Distributed Firewall
NSX Distributed Firewall with Advanced Threat Prevention	None required
NSX Gateway Firewall with Advanced Threat Prevention	None required
One of the following: ■ NSX Gateway Firewall with Threat Prevention ■ NSX Gateway Firewall	NSX Advanced Threat Prevention for Gateway Firewall
NSX Advanced with Advanced Threat Prevention	None required
NSX Enterprise Plus with Advanced Threat Prevention	None required

Software Requirements

You must also meet the following software requirements before you can start using the NSX Network Detection and Response feature.

- Install NSX 3.2 or later.

- Deploy NSX Application Platform. See *Deploying and Managing the VMware NSX Application Platform* document delivered with NSX 3.2 or later in the VMware NSX Documentation set.

  Note The versioning of the NSX Network Detection and Response feature that is hosted on the NSX Application Platform matches the NSX Application Platform version, and not the NSX product version number.

Important The NSX Network Detection and Response feature can function as designed only when your NSX environment is connected to the Internet. NSX Network Detection and Response is not supported in air-gapped environments when there is no outbound Internet access from the Kubernetes cluster pods and the NSX Unified Appliance.

Required Ports

Ensure that the required ports are open. Specifically, NSX Network Detection and Response requires the outbound TCP port 443 to be open. It uses this port to establish HTTPS connections to the NSX Advanced Threat Prevention cloud service and a limited set of other cloud services used to perform deeper threat analysis.

See the VMware Ports and Protocols webpage for other ports and protocols information.

Activate NSX Network Detection and Response

To activate the NSX Network Detection and Response feature, your NSX environment must meet the system requirements and the following prerequisites.

Prerequisites

- Ensure that the license, software, and port requirements are met. See System Requirements for NSX Network Detection and Response for details.
- You must be logged in using an Enterprise Admin account.

Procedure

1. From your browser, log in with the required privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
2. In the NSX Manager UI, select **System > NSX Application Platform**.
3. Navigate to the **Features** section, locate the NSX Network Detection and Response feature card, and click **Activate**.
4. In the NSX Network Detection and Response activation wizard, select one of the available cloud regions from which you can access the NSX Advanced Threat Prevention cloud service.

   The system uses the NSX Advanced Threat Prevention cloud service to perform deeper analysis of detected threat events, perform event correlation and visualization, and fetch periodic updates on those detected threat events. If you previously activated the VMware NSX® Malware Prevention feature, the cloud region selected for that feature is preselected and is used for the NSX Network Detection and Response feature.

5 Click **Run Prechecks**.

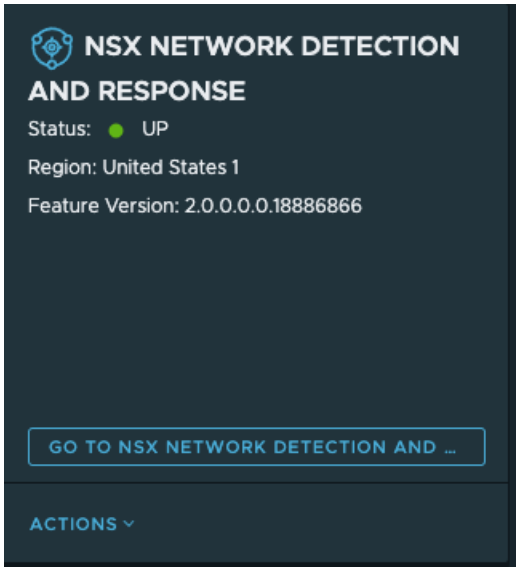
This precheck process can take some time as the activation wizard validates that the minimum license requirement is met. The wizard also performs the connectivity checks between the NSX Manager appliance and the NSX Advanced Threat Prevention cloud service. It also validates that the selected cloud region is reachable.

6 Click **Activate**.

This step can take some time to finish as the resources get allocated in the background.

Results

If the feature activation is successful, the NSX Network Detection and Response feature card displays the Status as `UP`, similar to the following image. It also displays information about the cloud region to which the feature is connected and the version of the feature that is being used.



See Troubleshooting NSX Network Detection and Response if you encounter any error during the NSX Network Detection and Response feature activation,

Working with the NSX Network Detection and Response Application

The VMware NSX® Network Detection and Response™ application offers a tightly integrated set of network detection and response capabilities for north-south and east-west security within your NSX environment. This feature is available beginning with NSX Intelligence version 3.2 and NSX version 3.2.

Prerequisites for Using the NSX Network Detection and Response Application

You must meet the following prerequisites before you can get started using the full functionalities of the NSX Network Detection and Response application.

- Get familiar with the main objective of NSX Network Detection and Response and understand its activation and usage workflow.

See Getting Started with NSX Network Detection and Response.

- Activate the NSX Network Detection and Response application on the NSX Application Platform.

  See Activate NSX Network Detection and Response

- Ensure you have an NSX role that is authorized to use the NSX Network Detection and Response feature.

  To access all the NSX Network Detection and Response functionalities, the user account you are using during your NSX Manager session must be assigned one of the following built-in roles. See Role-Based Access Control for more information.

  - Enterprise Admin
  - Security Admin
  - Security Operator
  - Auditor (Read-only access)

- Ensure that your NSX environment is connected to the Internet. NSX Network Detection and Response is not supported in air-gapped environments when there is no outbound Internet access from the Kubernetes cluster pods and the NSX Unified Appliance.

Terminology Used with the NSX Network Detection and Response Feature

Familiarize yourself with the following key terminologies that are used with the NSX Network Detection and Response feature.

Terminology	Definition
Campaign	A correlated set of incidents that affect one or more workloads over a period of time.
Event	Represents a security-relevant activity that has occurred in the monitored network. An event can involve multiple data flows (for example, TCP connections), but it represents a single type of activity occurring between a specific pair of IP addresses over a short period of time. Multiple events are automatically aggregated into incidents.
Incident	Represents a security-relevant activity that has occurred in the monitored network. An incident can consist of a single event or several events that have been automatically aggregated into an incident.
Infection	An incident that has been determined to be critical. Infections should be dealt with without delay.
Nuisance	An incident of low risk. This typically corresponds to potentially unwanted/risky activity that does not necessarily indicate a compromise or infection on the monitored network. Nuisances are tracked since they contribute to provide a more comprehensive network situational awareness.

Terminology	Definition
Event Impact Score	The overall impact score calculated for an event detected by the NSX Network Detection and Response feature. A score ranges from 0-100, with 100 being the most dangerous detection. The following levels of event impact are used. ■ Low: Impact 1-29 ■ Medium: Impact 30-69 ■ High: Impact 70-100
Watchlist	An incident that has been determined to be of medium risk. Such incidents, while indicating a potential risk, do not need immediate attention. They are kept under close watch in case new evidence appears that modifies their status. For example, an incident involving an inoperative command and control infrastructure is classified as watchlisted.

Getting Familiar with the NSX Network Detection and Response User Interface

The NSX Network Detection and Response user interface (UI) provides a single point of control for managing the threat events and correlated campaigns detected in your NSX environment, and viewing the generated reports about those threats.

Important To access the NSX Network Detection and Response user interface, you must first activate the NSX Network Detection and Response application on the NSX Application Platform. You also must activate one or more of the NSX features whose detection events the NSX Network Detection and Response application consumes. See Activate NSX Network Detection and Response.

Certain elements of the NSX Network Detection and Response user interface are visible only if you activate the element's corresponding feature or application that provides the events that the NSX Network Detection and Response application consumes.

Accessing the User Interface

If there are event reports or generated campaigns, you can access the NSX Network Detection and Response user interface (UI) using one of the following methods.

- Click the application launcher icon  in the upper-right corner of the NSX Manager UI and select **NSX Network Detection and Response**.
- Navigate to **Security > Security Overview** in the NSX Manager UI and in the **Threat Detection & Response > Campaigns** tab, click **Go to Campaigns**.
- If you activated the NSX Intelligence, navigate to **Security > Suspicious Traffic** in the NSX Manager UI. Expand the row for a detected suspicious event, click **Campaigns** or **Event Details**, if available. These links only appear if campaigns or event reports are available for the detected suspicious activity.
- If you activated the VMware NSX® Malware Prevention application, navigate to **Security > Malware Prevention** in the NSX Manager UI, expand the row for a reported malware, and click either the **Campaigns** or **Event Details,** if available. These links only appear if campaigns or event reports are available for the detected malware.

The following sections describe the common areas that you see as you navigate the NSX Network Detection and Response user interface. On the left side of the interface is the main navigation menu. At the top of almost every page are the display settings widgets. Data presented on the UI pages are displayed using the display settings that you have selected.

Navigating the Interface

You can use the main navigation menu on the left side of the browser page to access the corresponding top-level pages of the NSX Network Detection and Response UI. You can temporarily collapse this navigation menu, by clicking  in the upper-right corner of the menu panel. When you first see the NSX Network Detection and Response user interface, the **Dashboard** page is selected by default. The **Dashboard** page consists of widgets that provide an overview of multiple items being monitored. These widgets are described in more detail in Exploring the Dashboard Page.

To access another NSX Network Detection and Response interface page, click its corresponding tab on the main navigation menu on the left. Each tabbed page consists of several widgets that provide more information about the monitored areas. Topics available later in this guide provide details about each of these NSX Network Detection and Response UI pages.

Setting the Display Theme

You set the display theme used in your current NSX Network Detection and Response session using the display theme mode icon in the upper-right section of the interface. The icon that is displayed depends on the display theme that is currently in effect. To switch to a light-themed mode, click  . To switch to a dark-themed mode, click  .

Getting Assistance

To access the available NSX Network Detection and Response topics included in the *Using and Managing VMware NSX Intelligence* documentation, click  and then **Help**.

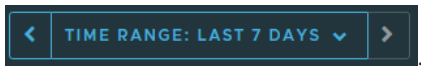
To see the status of your connection to the NSX Network Detection and Response cloud connector, click **Check connectivity status**. The cloud connector provides a secure tunnel of communication between your NSX Manager session and the NSX Advanced Threat Prevention cloud services.

If you encounter any connectivity issue that you cannot resolve using the information in the Troubleshooting section of this documentation, click **support ticket** and report the problem.

Accessing the Main NSX Manager User Interface

To return to the main NSX Manager user interface, click  in the upper-right corner, and select **NSX-T**.

Setting the Time Range

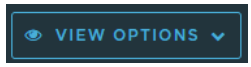
To specify the number of days of data to display in the NSX Network Detection and Response widgets, use the **Time Range** button 

. To navigate the date selection back and forward while keeping the selected range of dates constant, click  or  located on either side of the **TIME RANGE: LAST 7 DAYS** button. For example, assuming the default time range of 7-days, clicking the left arrow button once selects a range with the end date being 7 days ago.

You can define a more detailed time range using the **Time Range** pop-up window. Click **TIME RANGE: LAST 7 DAYS** button and select **Relative** (the default) or **Absolute** from the drop-down menu. In Relative mode, you select the number of days since the present date for which you want data displayed. The default is 7 days, the minimum is 1 day, and the maximum is 31 days. In Absolute mode, you enter the dates in **From** and **To** by selecting the dates from the calendar pop-up window. To save your selection, click **Apply**.

Using the View Options Button

All the date and time data that are displayed in the NSX Network Detection and Response interface uses the default UTC time zone, until you change it.

To change the time zone used for the displayed data, click 

located in the upper-right side of the interface and select the currently selected time zone. In the **Time Zone** pop-up window, click the drop-down menu and select a different time zone. To narrow the menu selection, start entering the name of a time zone in the search box. After you have selected the desired time zone, click **Apply**.

Managing the Widgets

Each of the NSX Network Detection and Response UI pages consists of multiple widgets that display details about the detected threats and reports generated from analysis of those threats.

You can manage the widgets using the following information.

- To reload the data displayed in a widget, click  at the top right corner of the widget.
- You can minimize a widget by clicking  or maximize it by clicking  next to the widget title.
- To focus further into the data displayed in some widgets, click the  icon.
- To view the data in XML/JSON format that is available in some of the widgets, click </>.
- Some widgets have contextual help displayed in a pop-up window. To access the help information, click  . In some contextual help pop-up windows, you can click the **here** link for more documentation about the data shown on the widget.

Exploring the Dashboard Page

The **Dashboard** page is where you start when reaching the NSX Network Detection and Response user interface.

The page provides a general overview of the active campaigns in your network, detected incidents and threats, and most recent observed threat events in your NSX environment.

The page consists of several widgets that can be managed using the information in Getting Familiar with the NSX Network Detection and Response User Interface.

You can drill down into the details of individual aspects of the interface and view detailed information. Some of this detailed information is presented directly in the widget. Other information is displayed on linked pages elsewhere in the NSX Network Detection and Response UI.

Active Campaigns in My Network

The **Active campaigns in my network** widget provides an overview of the campaigns that the NSX Network Detection and Response application identified and that are currently active in your network, surfacing the most critical campaigns for your immediate action.

The widget displays statistics for All Active Campaigns, Open High Impact Campaigns, In-Progress High Impact Campaigns, and Hosts Affected.

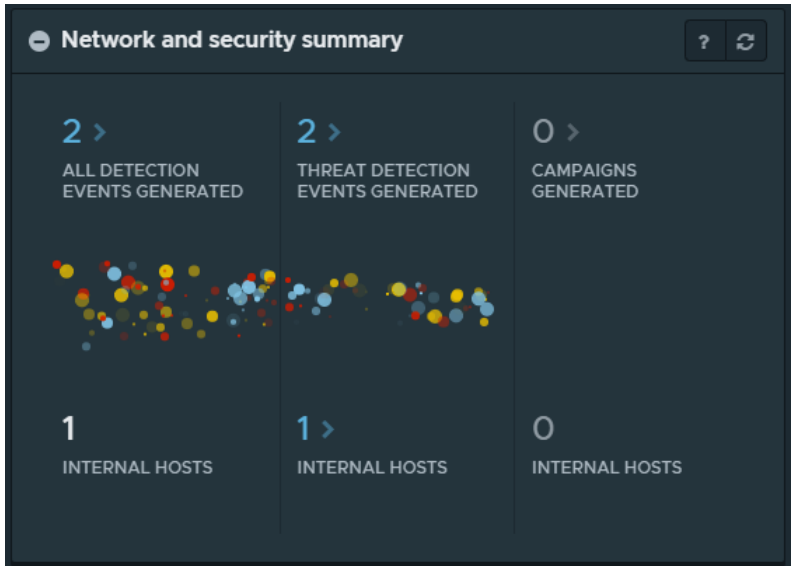
To see more details about these campaigns, click **Go to campaigns overview**, located in the bottom-left corner of the widget and you see the **Campaigns** page. See Managing the Campaigns Page for details.

Network and Security Summary

The **Network and security summary** widget shows how the NSX Network Detection and Response processes and analyzes network traffic flow data.

The widget shows the processing pipeline used for analyzing all events (including informational events), detecting threat events (only important events), and generating campaigns.

The widget has segments that indicate the different stages of processing the system performs on incoming data. As shown in the following image, the processing starts with All Detection Events Generated and proceeds until it reaches Campaigns Generated.

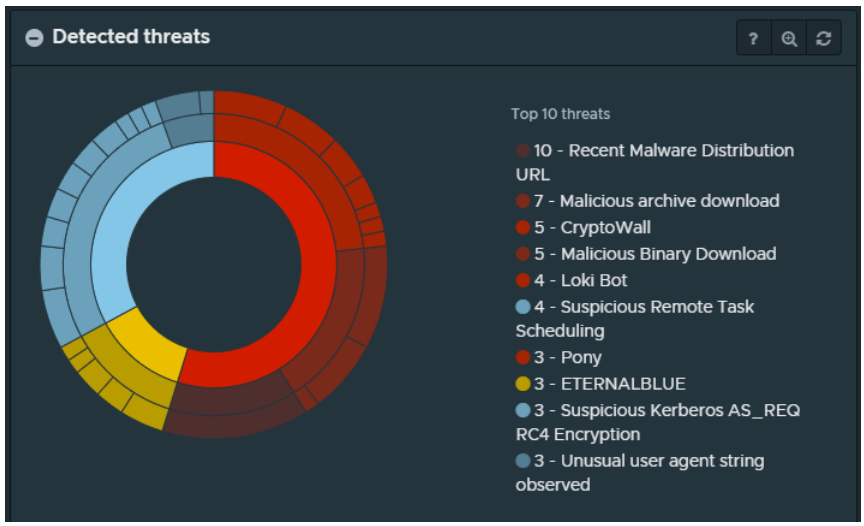


- When you click the count link for the All Detection Events Generated segment, you are taken to the **Events** page, which is filtered to display all the detection events, including unimportant info detection events. See Working with the Events Page for more details.
- Clicking the count link for the Threat Detection Events Generated segment takes you to the **Events** page, which has been filtered to display the list of threat detection events only. Clicking the count link below the Threat Detection Events Generated segment takes you to the **Hosts** page. See Working with the Hosts Page.
- When you click the count link for the Campaigns Generated segment, you are taken to the **Campaigns** page, which shows the cards for the detected campaigns. See Managing the Campaigns Page for details.

Detected Threats

The **Detected threats** widget provides a graphical overview of the different kinds of threats that the NSX Network Detection and Response application has detected in the network.

The threat information is displayed in a layered circle, similar to the following image.



The divisions of the circles represent the number of hosts affected by the displayed incident types. Moving toward the outer circles provides a finer granularity and more specific information.

- The innermost ring displays the three different types of incidents.

Incident Type	Description
Infections	These are incidents that the NSX Network Detection and Response application determined to be critical. These incidents have been given an impact score of 70 or higher and are displayed in red.
Watchlist	These are incidents that the NSX Network Detection and Response application determined to be of medium risk. Such incidents, while indicating a potential risk, might not need immediate attention. They are kept under close watch in case new evidence modifies their status. These incidents are assigned an impact score anywhere from 30–69 and they are displayed in yellow.
Nuisances	These are incidents that are considered low or no risk. This typically corresponds to potentially unwanted/risky activity that does not necessarily indicate a compromise or infection on the monitored network. These incidents have been given an impact score of lower than 30 and are displayed in blue.

- The middle ring displays the threat class together with the number of relevant incidents for each type of infection. Threat classes include command&control servers, malicious file downloads, crypto-miners, and many more.
- The outer ring represents the individual threat families detected in the network. Threat families include ransomware, malicious binary files, and so on.

When you point to the graph, the widget displays the threat name and a count of hosts where the NSX Network Detection and Response application observed the threat.

When you click an item in the graph, the view zooms in and displays more details about the selected information type. Clicking the item again zooms the view back.

If you click an incident type in the inner ring, the graph view zooms in and displays the matching incidents in the middle and outer ring. If you click a threat class in the middle ring, the graph view zooms in and displays the matching threat families. If you click the outer ring, the graph view zooms in and displays details about the selected threat.

The legend on the right side of the widget provides a count of the occurrences of the most frequent threats detected. When you point to an item in the legend, a pop-up window gives further information about the threat class, the number of incidents, and the number of affected hosts. Clicking the item zooms the graph view for the selected threat type and provides more contextual information.

Global Event Map

The **Global event map** widget provides a visual overview of geolocations of the aggregated events.

It marks the approximate location of the other hosts involved in the event detected by the NSX Network Detection and Response application. The marker color represents the event impact. The marker size represents the number of impacted hosts.

Events with no specific location are excluded from this map.

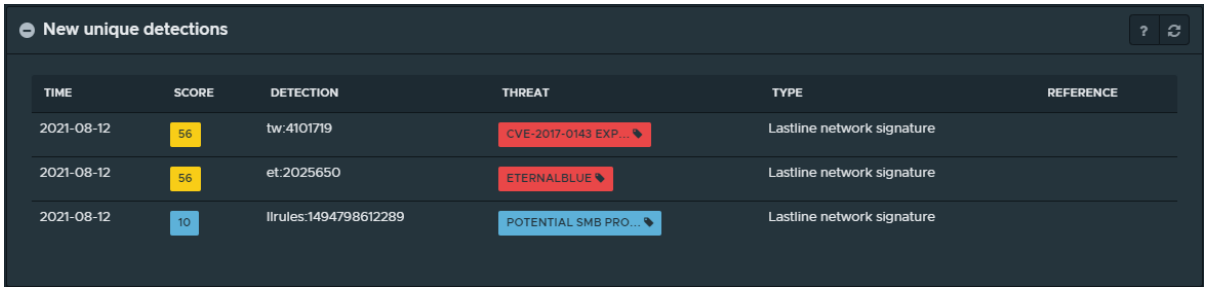
To learn more about the threats and hosts represented at that particular location, click a marker on the map.

In the **Location Details** pop-up window that displays, you can view the approximate location, the threats, and the destination hosts for the selected event. Click the ▼ icon next to each entry to apply filters to the list displayed in the **Events** page.

New Unique Detections

The **New unique detections** widget displays a list of events that the NSX Network Detection and Response application identified for the first time in your network.

The following image gives an example of the list displayed.



The list contains the time of the event, its impact score, the detection signature (which can be a URL with a malicious reputation or a specific rule), the threat, the event type, and a permanent link reference to the associated event.

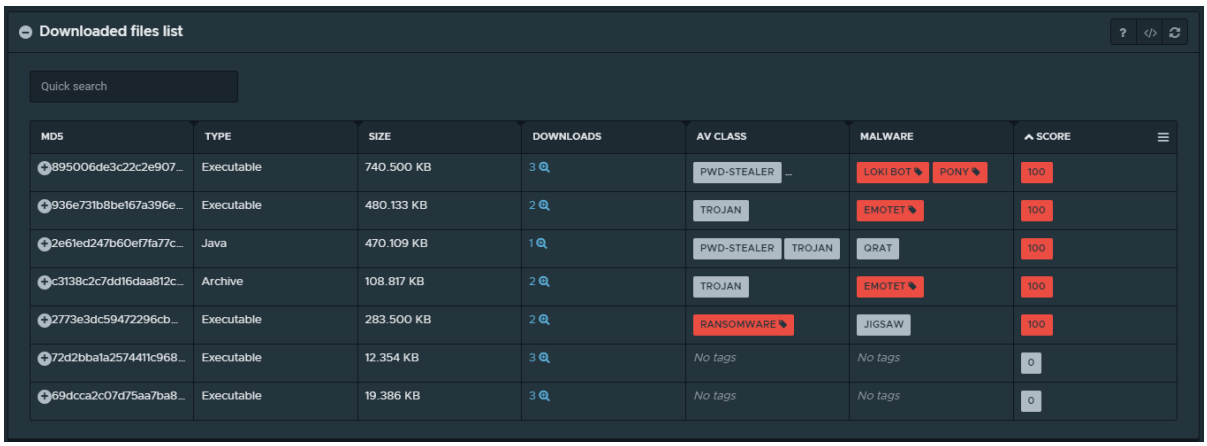
When you point to any of the rows in the list, more details are displayed about the detection event. Clicking the name of the threat displays a pop-up window with information about the threat type, severity, and details about the detected threat.

Important Triage and investigate these events since the NSX Network Detection and Response application detected these threats in your NSX environment for the first time.

Downloaded Files List

The **Downloaded files list** widget displays a list of distinct and unique files that the NSX Network Detection and Response application detected as having been downloaded by hosts in your network. This widget can only display data if the NSX Malware Prevention application is activated.

The following image shows an example of the **Downloaded files list** widget with data.



The **Quick search** text box in the upper-left corner of the list provides fast, as-you-enter search capability. It filters the rows in the list and displays only those rows that have text, in any column, that matches the query string that you entered in the search text box.

To customize the columns displayed in the list, click the ☰ icon located in the upper-right corner of the list.

You can customize the number of rows to be displayed. The default is 20 entries. Use the ‹ and › icons to navigate through multiple pages.

Each row is a summary of a downloaded file. Click the ⊕ icon or anywhere on an entry row to access a detailed view of the downloaded file.

The list is sorted by score and includes the following columns.

Column Name	Description
MD5	The MD5 hash of the downloaded file.
Type	The high-level file type of the downloaded file. Supported types are currently: ■ Archive – Archive formats such as ZIP or RAR ■ Document – Includes other types of Office documents ■ Executable – Binary application formats, such as Windows Portable Executable ■ Java – Java application or applet ■ Media – Macromedia (Adobe) Flash file ■ Other – Other recognized file format ■ PDF – Portable Document Format files ■ Script – An executable script such as JavaScript, Python, and others ■ Unknown – Unknown file type

Column Name	Description
Size	Size in bytes of the downloaded file.
Downloads	Number of times that the file was downloaded by hosts in the network. The displayed number and provide a link to the detailed downloads page. The link passes an Analyst UUID filter that restricts the view to downloads of this specific file.
AV Class	A label defining the antivirus class of the downloaded file. If the label has a , you can click that icon for a description in a pop-up window.
Malware	A label defining the malware type of the downloaded file. If the label has a , you can click that icon for a description in a pop-up window.
Score	The score assigned to the downloaded file by the analysis indicates the critical level of the detected threat and ranges from 0-100: ■ Threats that are 70 or higher are considered to be critical. ■ Threats that are between 30-69 are considered to be medium-risk. ■ Threats that are between 1-29 are considered to be benign. For details about maliciousness core and risk estimate, see Analysis Report: Overview Tab. If the icon appears, it indicates the artifact has been blocked. The list is sorted by decreasing order (most critical threats at the top). Click to sort the list in increasing order (least critical threats at the top), then click to toggle back to the default.

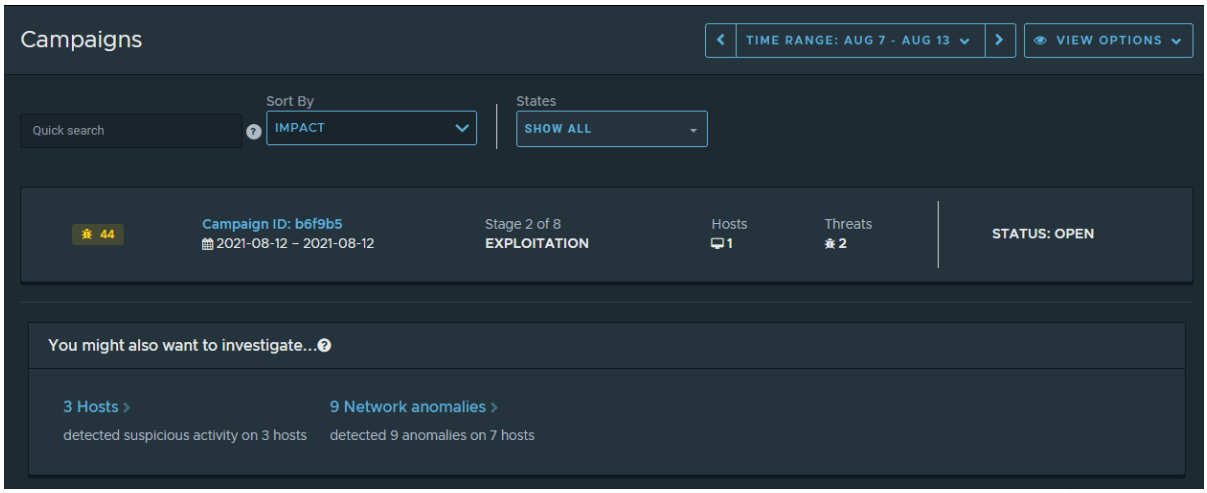
Managing the Campaigns Page

The **Campaigns** page provides an interface for monitoring the campaigns that the NSX Network Detection and Response application detected in your network.

The page consists of several widgets that can be managed using the information in Getting Familiar with the NSX Network Detection and Response User Interface.

If there are no detected campaigns, the `No campaigns found` message is displayed.

If there are detected campaigns, the page displays the corresponding campaign cards. The following image shows a sample **Campaigns** page with a card for the campaign detected during the selected time range. See Working with Campaign Cards.



At the bottom of the page, the **You might also want to investigate** widget is displayed. See About the Investigate Widget for more details.

Working with Campaign Cards

The **Campaigns** page displays campaign cards for any detected campaigns. A campaign card shows the calculated threat score, the campaign name (Campaign ID), the latest attack stage that the NSX Network Detection and Response application detected, the number of affected hosts, the number of different threats, and the campaign status.

Managing Campaign Cards

You can sort the campaign cards by clicking the **Sort By** drop-down menu and selecting from the list of criteria: **Impact** (the default), **Stage**, **Hosts**, **Threats**, **Newest**, or **Latest Activity**.

Select the campaign cards that you want displayed by clicking the **States** drop-down menu and selecting from **Show All** (the default), **Open**, **In Progress**, **Done**, or **Updated**. You can select more than one option. Clear a selection by clicking the option again.

To view all the available details about a campaign, click the **Campaign ID** link and the details about the campaign are displayed. See Understanding the Campaign Details Page.

Click anywhere on a campaign card and the **Campaign Summary** sidebar appears on the right side.

Understanding the Campaign Summary Sidebar

The **Campaign Summary** sidebar is displayed on the right side of the **Campaigns** page when you click anywhere on a campaign card.

The following describes what you see on the **Campaign Summary** sidebar.

Top section

At the top of the sidebar are the following items:

- The calculated threat score and the campaign name/ID (in long hash format) are displayed at the top.
- The **View Details** button, when clicked, gives you access to the **Campaign details** page. See Understanding the Campaign Details Page for more information.
- The number of hosts affected by the campaign is displayed.
- The number of threat types involved in the campaign is displayed.

Actions

The next section of the panel includes the following information.

- **Campaign Name/Campaign ID** – You can click the pencil icon and optionally edit the campaign name/ID.
- **State** – Select the triage status of the campaign from the drop-down menu. Select from **Open**, **In progress**, **Updated**, or **Done**.
- First Seen and Last Seen – Shows a linear graph with the timestamp from when the evidence was first and last seen. The Duration is displayed after the graph.

Attack Stages Seen

The **Attack Stages Seen** section displays the attack stages, highlighting the current campaign attack stages. Point to a highlighted activity (for example, **Exploitation**) to view a pop-up window with more information about the stage. See About Attack Stages for details.

Host Affected

The **Hosts Affected** section displays the hosts that are involved in the selected campaign. To view the **Host profile** page, click the IP address link. See Host Profile Page.

To see details about the hosts on the **Hosts** tab, click **View hosts**. See Campaign Details: Hosts Tab for more information.

Threats

The **Threats** section displays the current threats detected in the selected campaign. The color code indicates the severity of the threat: red for high severity, yellow for medium, and blue for low.

To view detailed information about the campaign on the **Campaign timeline** tab, click **View threats**. See Campaign Details: Timeline Tab for more information.

About the Investigate Widget

The **Investigate** widget displays the `You might also want to investigate ...` message and a list of custom-tailored facts and destinations prepared by the NSX Network Detection and Response application based on the activity in your network.

To explore interesting security information, follow the links provided in the list. The widget displays any of the following information depending on what the NSX Network Detection and Response application detected.

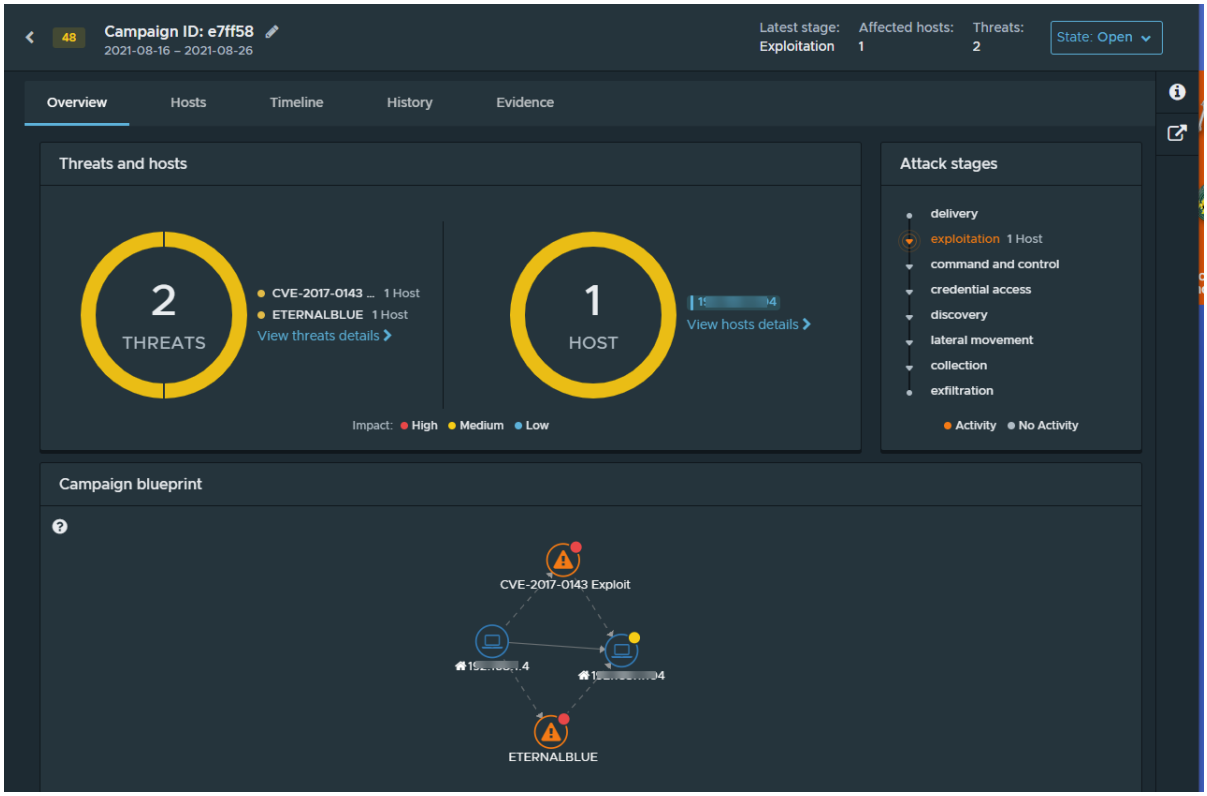
Detail Name	Description
Hosts	Reports suspicious activity on the hosts in your network. Click the link to go to the **Hosts** page.
Suspicious file downloads	Reports suspicious file downloads. Click the link to go to the **All** tab on the **Files downloaded** page.
Network anomalies	Reports `INFO` events that might need investigation. Click the link to go to the **Events** page.

Understanding the Campaign Details Page

The **Campaign Details** page on the NSX Network Detection and Response UI shows all of the available details for the campaign that you have currently selected in the **Campaigns** page.

You access this page by clicking a campaign's ID from the **Campaigns** page.

This page is divided into multiple tabs, as shown in the following image.



- **Overview** - Provides a summary and graphical blueprint of the campaign that the NSX Network Detection and Response application generated.
- **Hosts** - Provides a listing of the hosts affected by the campaign.
- **Timeline** - Displays the events included in the campaign in chronological order.
- **History** - Provides a textual history of the campaign.
- **Evidence** - Displays a list of the evidence detected for the currently selected campaign.

Across the top of the campaign details page is the data from the selected campaign card. It displays the calculated threat score, the campaign name (Campaign ID), the latest attack stage, the number of affected hosts, the number of different threats, and the state of the campaign.

To return to the **Campaigns** page, click the < icon in the upper-left corner of the page, next to the campaign threat score and campaign ID.

Campaign Details: Overview Tab

The **Overview** tab in the **Campaign Details** page displays a summary of the campaign and an interactive graphical blueprint.

The following information describes the three sections on this tab.

Campaign Threats and Hosts

The **Threats and hosts** section displays the **Threats** and **Hosts** widgets.

The **Threats** widget displays the current threats that the NSX Network Detection and Response application detected in the selected campaign. The severity of the threat is indicated by the color code: red for high, yellow for medium, and blue for low. Point to the name of the listed threats and a pop-up window displays the IP addresses of the affected hosts. Click **View threats details** and the **Timeline** tab displays detailed information about the campaign.

The **Hosts** widget displays the hosts affected by the selected campaign. The severity of the threat is indicated by the color code: red for high, yellow for medium, and blue for low.

Point to the IP address of an affected host and a pop-up window displays the names of the threats affecting the host. Click **View hosts details** and the **Hosts** tab displays detailed information about the hosts.

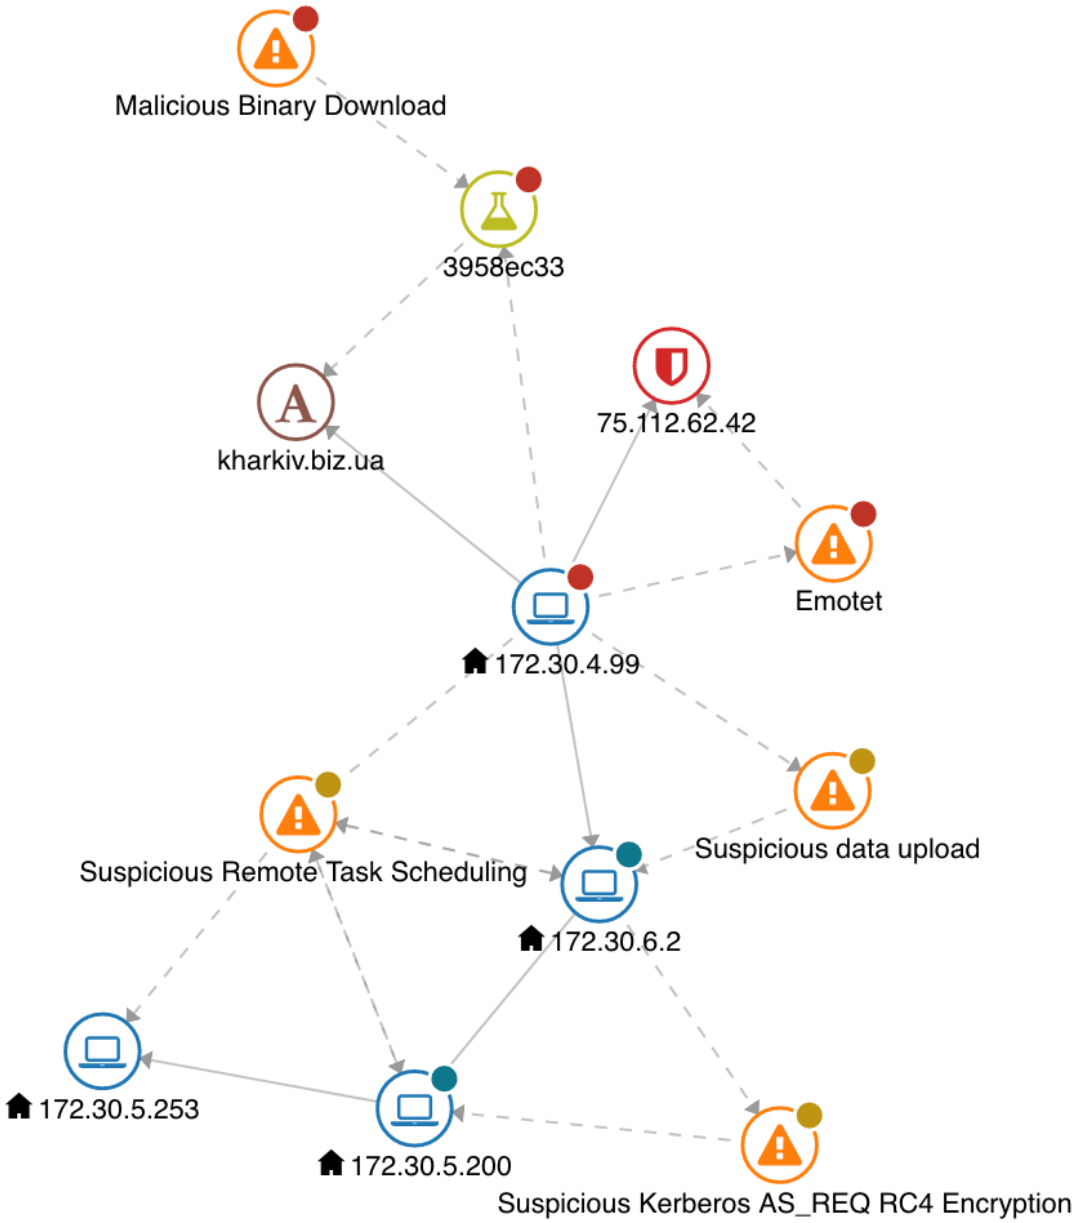
Campaign Attack Stages

The **Attack stages** widget displays the attack stages, highlighting the current campaign attack stages. Point to a highlighted activity and a pop-up window displays more information about the attack stage. See Campaign Properties for details about attack stages.

Campaign Blueprint

The **Campaign blueprint** widget provides an interactive graphical representation of the campaign. It displays the hosts involved in the campaign (both internal and external to your network), the threats that affected them, and additional information that completes the campaign description.

The following is an example of a blueprint graph.



This blueprint graph shows the following activities.

- A malicious binary file is downloaded to the host node with label 172.30.4.99. This activity is consistent with a user on that host opening an email (for example, visiting a URL or opening an attachment contained in that email).
- The host node with label 172.30.4.99 is connected to the hostname node with label kharkiv.biz.ua. The analysis report 3958ec33 shows that a download was made from the URL `http://kharkiv.biz.ua/hPpD/`. The analysis report also shows that what is downloaded is a PE executable application, 32-bit, Intel i386 file.

- The host node with label 172.30.4.99 is connected to an `Emotet command and control`. The server is the blocked entry 75.112.62.42.
- The host node with label 172.30.4.99 is connected to host node with label 172.30.6.2 with a suspicious data upload and to host nodes with labels 172.30.5.200 and 172.30.5.200 with a suspicious remote task scheduling, all activities associated with lateral movement.
- The host node with label 172.30.6.2 is connected to the host node with label 172.30.5.200 with a suspicious Kerberos encryption, an activity consistent with data exfiltration.

Node key

The following node types can appear in the blueprint graph.

Icon	Node type	Description
	Analysis report	This node type represents the results of detonating a sample (file or URL) in the NSX Network Detection and Response sandbox. ■ Analysis report nodes are labeled with a shortened version of the corresponding analysis task UUID. ■ The score range of the analysis run is expressed using the color-coded badge on the top-right of the node.
	Downloaded file	This node type represents a file that was downloaded in the network. ■ Downloaded file nodes are labeled with a shortened version of the corresponding file's SHA1 hash.
	Host	This node type represents a network device. ■ Host nodes are labeled with the IP address of the corresponding host. ■ The host node indicates whether a host is internal or external. Internal hosts display a  icon next to their IP address. The determination of whether a host is internal is based according to the private IP ranges configuration. ■ The maximum impact of incidents affecting the corresponding host is expressed using the color-coded badge on the top-right of the node.
	Info	This node type represents a detection of an info-level activity. This node only appears in the Network analysis blueprint graph. ■ An info event is created in the presence of activities or behaviors that are not necessarily malicious but provide additional, useful information. ■ The maximum impact of events detected for the corresponding threat is expressed using the color-coded badge on the top-right of the node.
	Threat	This node type represents a detection. ■ Threat nodes are labeled with the threat name associated with the detected event. ■ The maximum impact of events detected for the corresponding threat is expressed using the color-coded badge on the top-right of the node.

About Edges

The lines that connect the nodes are called edges.

A host node is connected to threat or analysis report nodes with a dotted line to indicate that the host corresponding to the host node was exposed to the threat represented by the threat or analysis report node.

Other connections are represented with a solid line to express that some activity (for example, a network connection, a DNS look-up, or a web request) put the entities corresponding to two nodes in relation.

Blueprint interaction

The blueprint graph is interactive: supporting item selection, moving nodes, and zooming in and out.

Node and edges can be selected by clicking on them: additional information about the selected item is found in the sidebar.

Hovering your mouse over a node colors the connecting edges, highlighting the interaction of that node.

Individual nodes can be dragged to new positions on the graph. The entire graph can be panned, effectively changing the point of view.

The graph can be zoomed in and out by scrolling the mouse wheel. More details are shown at higher zoom levels. In particular, the badge used with several node types to convey impact information is enriched with the actual impact score.

Campaign Sidebar

The **Campaign** sidebar is used to display information that is relative to one or more elements of the blueprint graph. By default, it is minimized.

- Click the icon to view node or edge information.
- Click the icon to view third-party tools.

To minimize the sidebar, click the icon.

Node or edge information

The node/edge information tab provides additional information about a selected node or edge in the blueprint graph. To select a node, click on its icon in the graph.

Node type	Information
Analysis report	Additional information about an analysis report. Report details: ■ Analysis reports – Displays the task UUID and score. Click the icon to view the analysis report in a new browser tab. ■ MD5 – File hash value. ■ SHA1 – File hash value. ■ Size – File size in bytes. ■ Category – The category the analyzed file belongs to. ■ Type – More detailed information about the file. Sightings details of the analyzed sample: ■ Number of downloads – The number of times the analyzed file was observed being downloaded. ■ Hosts – IP address of the hosts that downloaded the analyzed file. ■ URLs – The full URL of the downloaded file.
Downloaded file	Additional information about a downloaded file File details: ■ MD5 – File hash value. ■ SHA1 – File hash value. ■ Size – File size in bytes. ■ Category – The category the analyzed file belongs to. ■ Type – More detailed information about the file. Sightings details: ■ Number of downloads – The number of times the analyzed file was observed being downloaded. ■ Downloading hosts – IP address of the hosts that downloaded the analyzed file. ■ URLs – The full URL of the downloaded file. ■ Reports – Displays the report status, task UUID, and score. Click the icon to view the analysis report in a new browser tab.
Host	Additional information about a host. Host-level details: ■ IP address – Geo-located map or local network icon. ■ Hostnames – Domain name for the host. ■ Services – Any services detected on the host. Incidents involving the host: ■ Number of incidents – Count of all incidents. ■ Max impact – Indicates the maximum impact of all incidents. ■ Threats – A list of the detected events. A note indicates if the host is internal or external to the monitored network.

Node type	Information
HTTP request	Additional information about an HTTP request. URL details: ■ Download URLs – The observed URL(s) in the HTTP request. ■ Download IPs – The IP address(es) resolved for the HTTP request. Click the icon to view the request IP address in Network analysis. Request details ■ Number of requests – The number of times the HTTP request was observed. ■ Hosts – IP address of the hosts issuing the HTTP request. ■ Referers – The "referer" header values observed in the HTTP request. ■ User agents – User-agent values observed in the HTTP request.
Threat	Additional information about a threat Threat details: ■ Threat class – The name of the detected threat class. For example, command&control. ■ Threat – The name of the detected threat. For example, Loki Bot. ■ Severity – The calculated threat score. ■ Information – a description of the detected threat

When you click an edge, the following information is displayed about the connection:

- Source node – The source of the connection. This can be a node name, an IP address, a domain name, etc.
- Target node – The destination of the connection. This can be a node name, an IP address, a domain name, etc.

Under the Source node and Target node is the actual source or target of the connection. Click the icon to expand the source or target.

Third-party tools

The third-party tools tab links to external tools that may provide additional information about an entity selected in the graph. Currently, the tools supported are DomainTools and VirusTotal.

The following searches are supported:

- Selecting a host node allows you to search for the corresponding IP address on DomainTools and VirusTotal.
- Selecting a hostname node allows you to search for the corresponding domain name on DomainTools and VirusTotal.
- Selecting a downloaded file node allows you to search for the corresponding hash on VirusTotal.
- Selecting an HTTP request node allows you to search for the request's hostname on DomainTools and VirusTotal.

Campaign Details: Hosts Tab

The campaign **Hosts** tab in the **Campaign Details** page displays a list of hosts that the campaign affected.

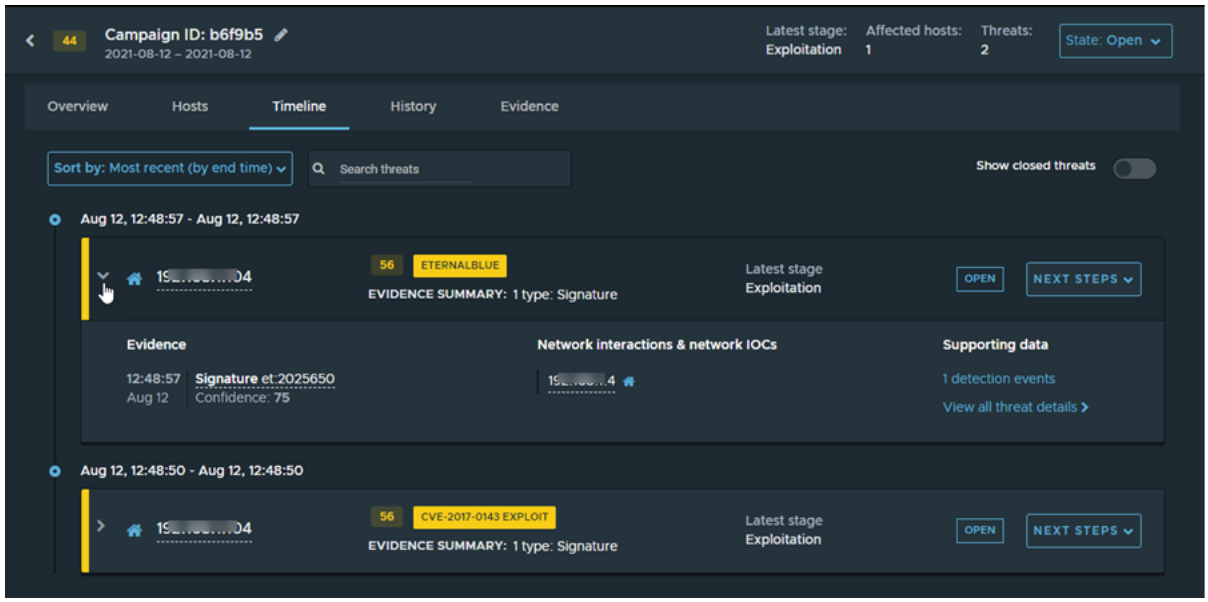
The columns provide the following information.

Column Name	Description
Hosts	The IP address of the host affected by the campaign. Click the IP address link and the **Host Summary** sidebar displays on the right side.
Threats	A list of all the threats that NSX Network Detection and Response detected on the host.
Attack Stages	The attack stages observed during the threat activity affecting that specific host.
Latest Activity	The timestamp of when an activity was last detected for that host.

Campaign Details: Timeline Tab

On the campaign's **Timeline** tab in the **Campaign Details** page, the threats detected by NSX Network Detection and Response are represented by threat cards.

A threat card displays the host that is connected to this threat, the calculated threat score, the threat name and class, the detection outcome (if available), the threat status, and other actions. To see its related evidence, expand the card by clicking the › icon, as shown in the following image. Click the ⌄ icon to collapse the Evidence section.



Sort the threat cards with the **Sort by** drop-down menu. Select from **Most recent** (the default), **Earliest**, **Highest impact**, and **Lowest impact**.

The **Search threats** text box above the list provides fast, as-you-enter search. It filters the rows in the list, displaying only those rows that have text, in any field, that matches the query string. Your query is matched against values across the following categories: impact, IP address, threat/malware, latest campaign phase, first seen, evidence, and other hosts, and, for mail messages, message information.

To filter the displayed threat cards by threat status, toggle the **Show closed threats** button. The default is to show all threats.

Threat Cards

The threat cards show all the threats associated with the selected campaign and their corresponding threat levels.

Each card displays the calculated threat impact, the threat name, the threat class, and if available, the detection outcome. It also shows the status of the threat: `OPEN` or `CLOSED`.

You can click **Next Steps** and select an action from the drop-down menu. Select **Close** to close the threat, **Open** to reopen a closed threat, or **Manage Alert** to create an alert management rule from the threat.

The **Evidence Summary** section contains an overview of the evidence and other data detected for the threat. Click the ⌄ icon (or almost anywhere else in the card) to expand the Evidence details section.

Evidence Details

The **Evidence** column displays the file downloads, signatures, and other categories along with a timestamp of when the evidence was seen.

The **Network interactions & network IOCs** column displays the IP address or domain name of external hosts. Click the IP address link to expand the **Network Interaction** sidebar.

The **Supporting data** column provides a link to the detected events, a link to the captured data, and a link to the threat details.

Campaign Details: History Tab

The campaign **History** tab in the **Campaign Details** page of the NSX Network Detection and Response UI displays a descriptive textual history of how the campaign was formed.

Each entry provides a Notice and a Description of the recorded campaign phases, along with the Notice Timestamp.

Campaign Details: Evidence Tab

The **Evidence** tab in the **Campaign Details** page of the NSX Network Detection and Response UI displays a list of the evidence detected for the currently selected campaign.

Each row is a summary of the evidence for the campaign. Click ⊕ (or anywhere on an entry row) to expand the row to view the Signature evidence information.

The evidence list includes the following columns.

Evidence Columns	Description
IP Address	The IP address of the host that is the source of the threat.
First Seen	Timestamp showing the start time of campaign.
Last Seen	Timestamp showing the most recent activity of the campaign.
Threat	Name of the detected security risk.
Threat Class	Name of the detected security risk class.

Evidence Columns	Description
Impact	The impact value indicates the critical level of the detected threat and ranges from 1-100: ■ Threats that are 70 or above are considered to be critical. ■ Threats that are between 30-69 are considered to be medium-risk. ■ Threats that are between 1-29 are considered to be benign. If the [block icon] appears, it indicates the artifact has been blocked.
Evidence	The derived value of the evidence for the campaign. See About Evidence for details.
Subject	Additional information from the campaign. This may be an IP address or an HTTP response code, or some other data.
Reference	Click the link to access the **Network event details** page. The link opens in a new browser tab. See Event Profile Page for details.
Incident ID	A permalink to a correlated incident. The link will open in a new browser tab. See Managing the Incidents Page.

Click the ☰ icon to change which columns to display. The default is to display all available columns.

When you click ⊕ (or anywhere on an evidence row), the following information is shown.

Information Name	Description
Threat	Name of the detected security risk.
Threat class	Name of the detected security risk class.
Impact	The impact score of the campaign.
Detector	If present, displays the NSX Network Detection and Response module that identified the threat. Click the link to view the **Detector** pop-up window.
View network detection	If present, displays the NSX Network Detection and Response module that identified the threat. Click the link to view the Detector pop-up window.
View Incident	Click the link to access the Network event details page. The link opens in a new browser tab. See Event Profile Page.
First seen	Timestamp showing the start time of campaign.
Last seen	Timestamp showing the most recent activity of the campaign.
Severity	An estimate of how critical the detected threat is. For example, a connection to a command and control server is typically considered high severity as the connection is potentially damaging.
Confidence	Indicates the probability that the detected individual threat is indeed malicious. As the system uses advanced heuristics to detect unknown threats, in some cases, the detected threat may have a lower confidence value if the volume of information available for that specific threat is limited.

Campaign Properties

A campaign detected by the NSX Network Detection and Response application is characterized by multiple properties.

The following are the campaign properties and their definitions.

Property Name	Description
Name	A campaign ID that uniquely identifies the campaign.
Hosts	The hosts that are affected by the campaign.
Threat	The threats that have been detected for the campaign.
Attack Stages	The phases in the life cycle of the adversary corresponding to the detected activities. See About Attack Stages for details.
Duration	The time interval during which the activities associated with a campaign have been observed.

About Attack Stages

Attack stages are the phases in the life cycle of an adversary that corresponds to the activities detected by the NSX Network Detection and Response application.

An adversary model describes the actions that an adversary may take to compromise and operate within an enterprise network. The NSX Network Detection and Response application uses MITRE's Adversarial Tactics, Techniques, and Common Knowledge (ATT&CK™) model to describe adversary behaviors. In this model, the techniques that an adversary could use are grouped into a number of tactic categories, which correspond to different stages in the attack lifecycle.

In the system, the activity associated with each detected event might be associated to a specific attack stage and might provide an indication of the campaign progress along its lifecycle. (Activities encountered at different attack phases might not be associated to a specific attack stage.) Currently, the following attack stages are used.

Attack Stage Name	Description
Delivery	The stage where attackers send the payload to the target. Common delivery mechanisms include remote exploits, drive-by-download webpages, and malicious USB or other removable drives.
Exploitation	The stage where the attacker's payload is deployed in the target network. Consequently, one or more devices in the target network are compromised and under the attacker's control.
Command and Control	The stage where attackers communicate with systems under their control within the target network, effectively obtaining "hands on the keyboard" remote access to these systems.
Credential Access	The stage where attackers gain access or control over system, domain, or service credentials used within the target environment. Typically, attackers attempt to obtain legitimate credentials from user and administrator accounts in order to impersonate them, or to create new accounts.
Discovery	The stage where attackers attempt to find more information about the target environment. Attackers often attempt to identify additional devices in the network, which they can use for their objectives.
Lateral Movement	The stage where attackers move across the target network by gaining access and control of remote systems.
Collection	The stage where attackers identify and gather information from a target network prior to exfiltration.
Exfiltration	The stage where attackers remove files and information from a target network.

About Correlation Rules

In general, incidents are grouped into a campaign when there is evidence that indicates that the corresponding malicious activities or attacks are related.

Since these correlation rules are running in the NSX Advanced Threat Prevention cloud service, they can be improved or extended independent of the NSX release cycles. Additionally, the list of correlation rules or the specific behavior of a rule can change over time.

The following are the currently supported correlation rules.

Anomaly Event

This rule correlates the detection events from the NSX Suspicious Traffic feature with higher-impact infection-type events. For example, an anomaly event from the NSX Suspicious Traffic feature coincides with a high-impact network event for the same hosts.

Exfiltration

This rule correlates exfiltration events that are preceded by infection-type events. For example, a command and control network event is followed by a network event that we know is exfiltrating data.

File Transfer (Hash-Based)

This rule correlates malicious file transfers. For example, if the same malicious file is downloaded to a number of hosts in the network, the rule will correlate all of these transfers into an intrusion. The similarity of malicious file transfers is determined based on the transferred file's SHA-1 hash.

File Transfer (Analysis Tag-based)

This rule correlates malicious file transfers. For example, if the same malicious file is downloaded to a number of hosts in the network, the rule will correlate all of these transfers into an intrusion. The similarity of malicious file transfers is determined based on the tags associated to the analysis tasks of the files.

Vulnerability Scan

This rule correlates different types of network events that all potentially indicate a vulnerability scan. For example, multiple outbound infection-type or NTA events are observed from a single host towards one or multiple internal destination hosts.

Wave

This group of rules identifies attack "waves," in which the same attack (that is, incidents for the same threat) is observed on multiple hosts across the network within a certain time window.

This group of rules is useful to identify hosts in the network that have become part of the same command and control infrastructure or have been exposed to the same attack vector (for example, drive-by attack or malware distribution attack). As a result, these rules are restricted to threats of class command and control, drive-by, malware distribution, sinkhole, fake AV, and crypto mining.

The rules in this group trigger in the following cases.

- There are network signature events where the threat class is command and control, affecting multiple hosts.

- There are network signature events where the threat class is malware distribution, affecting multiple hosts.
- There are network signature events where the threat class is drive-by and the entry (IP address or hostname) where the detections occurred are the same, affecting multiple hosts.
- There are malicious reputation events for the same entry (IP address or hostname) and the threat class is command and control, affecting multiple hosts.
- There are malicious reputation events for the same entry (IP address or hostname) and the threat class is malware distribution, affecting multiple hosts.

In this case, the correlation window is set to three days. Therefore, two incidents for the same threat affecting different hosts are considered related if they occur within this limited time range.

Note These rules can create campaigns comprising only of one host and one incident.

Confirmed Drive-by

This group of rules identifies campaigns where an internal host is exposed to a successful drive-by attack. A drive-by attack on a host is considered successful if it is followed by command and control, malware download, sinkhole, or fake AV activity. The rules in this group trigger in the following cases.

- Drive-by closely followed by malware download activity: In this case, the correlation window is 10 minutes, as we expect the download to be immediately caused by a successful browser exploit.
- Drive-by closely followed by fake AV activity: In this case, the correlation window is 10 minutes, as we expect the fake AV activity to immediately follow a drive-by exploit.
- Drive-by followed by command and control activity: In this case, the correlation window is four hours, as the command and control channel might take some time to be set up.
- Drive-by followed by sinkhole activity: In this case, the correlation window is four hours, as the activity towards a sinkholed malicious server over a command and control channel might take some time to be set up.

Note These rules can create campaigns comprising only of one host.

Confirmed File Download

This group of rules identifies campaigns where a malicious file is downloaded and successfully executed on a host. A downloaded file is considered to have successfully executed on a host if, shortly following the download, there are network events for activities that match the activity observed during the file analysis.

In particular, the file analysis can provide two more pieces of information to characterize the activity observed during the analysis.

Malware information

If the file behavior matches the behavior of a well-known threat, the malware name becomes available.

Network IoC information

If during the analysis the sample generates network traffic matching network signatures or threat intelligence, indicators for the traffic are made available. That is, information about malicious reputations and network signature matches are provided.

The rules in this group trigger in the following two cases, depending on the type of information derived from the file analysis.

- Malware-based case
  - A file is downloaded on a host.
  - The file analysis attributes a specific threat to the file (for example, Emotet malware).
  - At a later time, a network event for the same threat (that is, Emotet) is detected for the host that downloaded the file.
- Network IoC-based case
  - A file is downloaded on a host.
  - The file analysis identifies network IoC for the file.
  - At a later time, the host that downloaded the file attempts to contact an IP address or hostname included in the malicious reputation IoC extracted for the file and this traffic matches a network signature.

The NSX Network Detection and Response application sets the correlation window in this case to three days.

Note This rule can create campaigns comprising only of one host.

Lateral Movement

This group of rules identifies campaigns where attackers have established a "beachhead" in the network by compromising some hosts and then attempt to move laterally within the network to compromise additional hosts.

This group comprises two rules, each of which detects a separate step of the lateral movement campaign.

Outgoing lateral movement

This rule correlates outgoing lateral movement activity from a host in the home network and infections on that host that happened before the lateral movement detections (but within the correlation window).

Incoming lateral movement

This rule correlates incoming lateral movement activity towards a host in the configured home network and activity commonly observed after an initial compromise (Command&Control, probing and credential harvesting) that occurred on the same host after the lateral movement detections.

Notice that these rules will only trigger for hosts within the home network, that is the campaign is created only if both source and destination hosts of the lateral movement activities belong to the home network. If the home network is not configured, the system uses RFC1918 ranges by default.

INFO Events Promotion

The NSX Network Detection and Response application detects several activities in a protected network that might be interesting to an analyst, but are probably not malicious. These detections generate INFO events, which can be viewed by setting an appropriate value of the "event outcome" filter.

The NSX Network Detection and Response application does not consider INFO events for correlation purposes.

A challenge with these detections is that the same INFO event activity can be normal or highly suspicious, depending on the network in which the NSX Network Detection and Response application detected it. For example, the use of the remote desktop protocol (RDP) can be normal in an environment where this tool is used for legitimate administrative purposes, but can otherwise be a highly suspicious indication that an attacker might be attempting to remote-control a victim host.

Anomaly detection logic is able to determine when certain kinds of INFO detections are unusual for the monitored network and for the specific source hosts and destination hosts involved. When the system determines that an INFO detection is unusual, the event is promoted to "detection" mode and, as a consequence, is displayed among regular events. This scenario is relevant in the context of correlation rules for lateral movement, as the detection of lateral movement activity often result in the creation of INFO events.

Home Network

The home network configuration has the following effect on campaign correlation rules.

- All campaign correlation rules ignore events that happened on hosts outside of the home network.
- If no home network is configured, the system defaults to the RFC1918 ranges.

The home network is configured in **Security > General Settings > Private IP Ranges**.

Host Silencing

The host silencing configuration has the following effect on campaign correlation rules.

- If host silencing is configured, all campaign correlation rules ignore events that happened on silenced hosts.

- If no host silencing is configured, all source hosts detected in an event are considered valid for correlation.

To ensure that host silencing does not mistakenly include hosts whose activity should be included in campaigns, you must verify your host silencing configuration.

About Evidence

The NSX Network Detection and Response application reports on the actions observed while analyzing an event, incident, or campaign.

The evidence contains the following information.

Basic Detection Evidence: Network

Evidence type REPUTATION

Indicates that network traffic was detected to an IP or domain that is associated with a known threat.

A SUBJECT field and an IP address or domain are displayed. For example: reputation: `evil.com` (reference event), `6.6.6.6` (reference event), or `bad.org` (reference event).

These bad domains and IP addresses are typically blocked. Additional reputation information is displayed if available.

IP addresses may be annotated with a location (country flag).

Evidence type SIGNATURE

Indicates that network traffic was detected that matches a network signature for a known threat.

A Detector field that is the name/unique identifier of the signature that matched is displayed. For example, `Detector: et:2014612` or `Detector: llrules:1490720342088`.

Evidence type ANOMALY

Similar to SIGNATURE with the difference that detection is based on a heuristic that detected something anomalous. For example, `Anomaly: anomaly:download_smb`.

Evidence type FILE DOWNLOAD

A malicious or suspicious file was downloaded.

A `task_uuid`, the identifier of an analysis (detonation in sandbox), and the `severity`, the score of that analysis, is displayed. For example, `File download: a7ed621`.

The following lists additional optional information from the reference event.

- The URL the file was downloaded from
- The file type (typically executable)
- The filename

Evidence type UNUSUAL_PORT

Indicates that a TCP or UDP port is being used that is an uncommon one and corresponds to what is expected of this specific threat.

The IP address or domain involved in the traffic that used the unusual port is displayed in the SUBJECT field.

Evidence type URL_PATH_MATCH

Similar to unusual port, with the difference that detection is based on a URL path. For example, `http://evil.com/evil/path?evil=threat`, the detection is triggered by the `/evil/path` portion of the URL.

Evidence type DGA

DGA stands for "Domain Generation Algorithm", an approach used by some malware, where instead of using a small number of domains for Command and Control, the malware includes an algorithm that generates thousands of new random-looking domains each day. It then tries to contact each of them. To control their malware, the hacker just registers one or a few of these domains. The use of DGA is very visible on the network due to resolution attempts of many such domains.

The DGA evidence is currently used in addition to regular reputation evidence, when multiple bad domains from a DGA algorithm being resolved is detected.

Evidence from Correlation of Multiple Events

Evidence from correlation of multiple events

The following evidence types are created in cases when the combination of multiple network events on a host increases confidence that a threat has been correctly detected. The evidence types may be, for example, the same malicious reputation entry being contacted or the same network signature being triggered.

For each of these cases, the threat may be tagged as follows.

- `Repeated`: The specific threat was seen three or more times.
- `Periodic`: The specific threat was also seen occurring at regular intervals.

A label is shown on the corresponding reputation/signature evidence.

In the example of REPUTATION evidence, if repeated and periodic evidence for `bad.org` are detected, a REPEATED or PERIODIC tag is displayed.

Evidence type CONFIRMED_EXECUTION

This is associated with threats, such as MALICIOUS FILE DOWNLOAD. It means that network behavior is detected from the host that downloaded the file that confirms that the downloaded file was actually executed.

That is:

- A malicious file was downloaded to host `1.2.3.4`.

- When executed in a sandbox, this file contacted evil host `evil.com`.
- Shortly afterwards, command and control traffic from host `1.2.3.4` to `evil.com` is observed, confirming that the malicious file was executed.

The linked reference event is to where file was downloaded.

Additional evidence can provide confirmation of the threat, such as the following information about the file.

- Task UUID
- Score
- Filename
- URL it was downloaded from

Evidence type CONFIRMED_C&C

Similar to CONFIRMED_EXECUTION, this evidence is added to the command&control detection for the specified threat because the host previously downloaded a file for that threat.

Evidence type CONFIRMED_DRIVE_BY

This is added in situations a drive-by attack was detected followed by some indication that attack was successful. For example:

- Host `1.2.3.4` seems to be the victim of a drive-by attack.
- Shortly afterwards, host `1.2.3.4` either:
  - Downloaded a malicious file
  - Performed command&control traffic

This evidence is added to reference event of the initial drive-by event.

Evidence type DRIVEBY_CONFIRMATION

Similar to CONFIRMED_DRIVEBY evidence, this evidence is added as a reference event to the malicious file download or command&control detections that happened shortly after a drive-by attack.

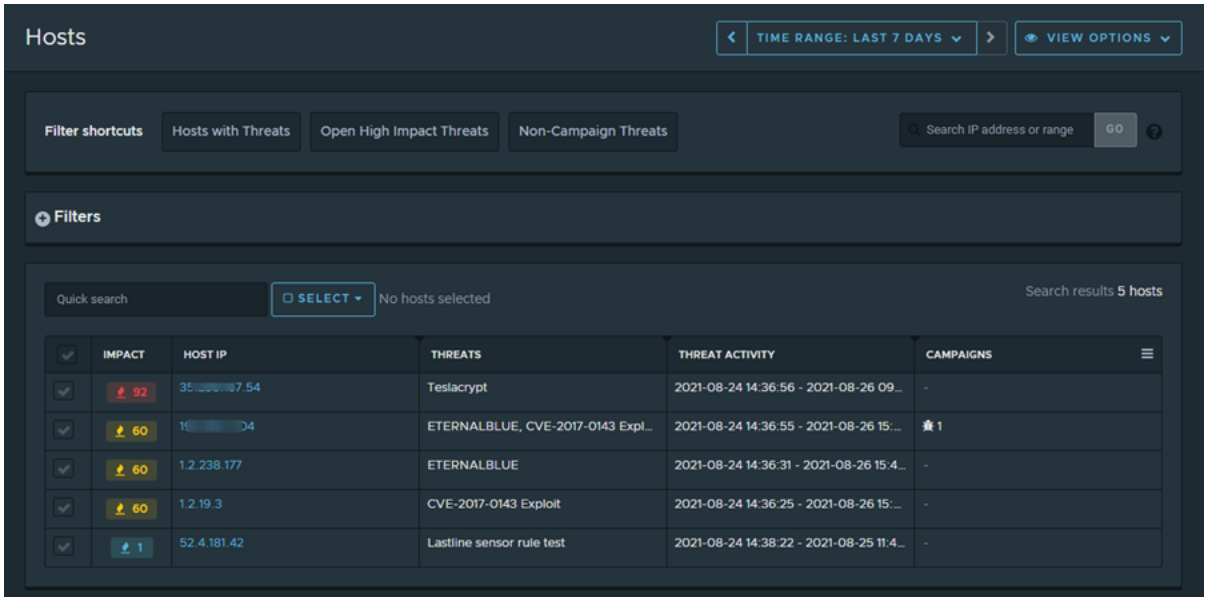
Working with the Hosts Page

The **Hosts** page displays a list of the monitored hosts in your NSX network.

The page consists of several widgets that can be managed using the information in Getting Familiar with the NSX Network Detection and Response User Interface.

You can quickly customize the selection of hosts using the filter shortcuts. You can also select your own filters. Use these filters to customize the Hosts list that displays on the **Hosts** page.

The following image shows a sample **Hosts** page.



Filter Shortcuts

To limit the data in the hosts list displayed on the **Hosts** page of the NSX Network Detection and Response user interface, use the filter shortcuts.

To select one of the following filter shortcuts, click the corresponding button displayed in the UI.

Shortcut Name	Description
Hosts with Threats	List all hosts on the home network with detected threats.
Open High Impact Threats	List all hosts on the home network with high impact threats that are open.
Non-Campaign Threats	List all hosts on the home network with threats that are not part of an campaign.

Alternatively, you can also limit the data displayed by entering a valid IPV4 IP address, range of IP addresses, or CIDR block in the search text box on right side of the **Filter shortcuts** widget and clicking **GO**.

Use Filters on the Host Page

The NSX Network Detection and Response application provides a filtering mechanism that allows you to focus on specific host information that is of interest to you. The use of filters is optional.

Procedure

1. From the **Hosts** page, click ⊕ to expand the **Filters** widget.
2. Click anywhere in the **Filter on** text box and select an item from the drop-down menu.

   You can select from the following available filters. To further narrow the focus of the displayed information, you can combine multiple filters.

Filter Name	Description
Campaign UUID	Restrict the displayed entries by the Campaign UUID. This is a 32-character hexadecimal string, for example, `7dabc0fc9b3f478a850e1089a923df3a`. Alternatively, enter the string `null` to select records that do not belong to any campaign.
Home network	Restrict the displayed entries by the Home network setting. Select **Home network only** or **Unidentified networks only** from the drop-down menu.
Host IP	Restrict the displayed entries to a specific source IP address, IP address range, or CIDR block. Enter the value in the text box.
Hosts with threats	Restrict the displayed entries by the Hosts with threats status. Select **Only hosts with threats** or **All hosts** from the drop-down menu.
Priority	Restrict the displayed entries by the Priority status. Select **Infections**, **Watchlist**, or **Nuisances** from the drop-down menu.
Read	Restrict the displayed entries by their Read status. Select **Read** or **Unread** from the drop-down menu.
Status	Restrict the displayed entries by their status. Select **Closed** or **Open** from the drop-down menu.
Threat	Restrict the displayed entries by a specific threat. Select a threat from the drop-down menu. The menu is prepopulated with a list of cataloged threats. Use the search function at the top of the menu to quickly find a threat name.
Threat class	Restrict the displayed entries to a specific class of threats. Select the threat class from the drop-down menu. The menu is prepopulated with a catalog of classes, some of which are listed below. Use the search function at the top of the menu to quickly find a class name. ■ **adware**: Malware that displays or downloads advertisements to an infected computer. ■ **click-fraud**: Click-fraud targets pay per click online advertising. ■ **command & control**: An infected machine belongs to a botnet and the machine can be remotely controlled by an attacker. ■ **drive-by**: An attacker attempted to exploit a vulnerability on the machine in order to install additional malware on the target system. ■ **exploit toolkit**: Detection of an exploit toolkit that attempted a drive-by download attack ■ **fake-av**: Fake antivirus software or other kinds of rogue security software designed to trick or mislead your users. ■ **inactive C&C**: The command & control server for this specific botnet is inactive. ■ **Malicious File Download, Malware Distribution, and malware download**: The IP address or domain hosts malicious executables. ■ **sinkhole**: A sinkhole is operated by a legitimate organization, so it does not pose a threat. However, hosts that try to contact such a host may be infected. ■ **spyware**: Malware that attempts to steal sensitive information. ■ **suspicious-dns**: Suspicious DNS domains are domains that are contacted by malware running on infected machines. Our proprietary techniques were able to proactively identify these domains as malicious. ■ **unknown**: An unknown security risk was detected.

3 To apply the selected filters, click **Apply**.

4 (Optional) To delete an individual filter, click the **Remove –** button next to its entry. To delete all the selected filters, click the ✖ icon located on the right side of the **Filter** widget.

The **Filters** widget collapses when you delete all the selected filters.

Hosts List

The bottom portion of the **Hosts** page of the NSX Network Detection and Response UI displays a list of hosts that meet the criteria of the selected filters. If no filters have been selected, all hosts in your network are displayed.

Search

The **Quick search** text box in the upper-left section of the list widget provides a quick as-you-enter search feature. The system filters the rows in the list, displaying only those rows that have text, in any column, that matches the query string.

Note If the list is long, the **Quick search** only scans the first 1,000 entries and can return incomplete results. The total number of search results returned is displayed in the upper-right corner of the list widget.

Selection

Use the **SELECT** drop-down menu for a fine-tuned selection. The available selection options are **All visible**, **All pages**, or **Clear selection**. To select all visible hosts, you can also click ☑ in the row for column names.

Host List

You can customize the number of rows that are displayed in the Host list. The default is 20 entries. To navigate through multiple pages, use the ❮ and ❯.

Each row provides an information summary for a host. To select a host row, click the ☑ icon. To access more information about a host, click anywhere in an entry row and the **Host Summary** sidebar panel displays. See Host Summary Sidebar for details.

The list includes the following columns.

Column Name	Description
IMPACT	Active threats on the host are denoted with the icon. The impact value indicates the critical level of the detected threat and ranges from 1–100: ■ A threat value that is 70 or higher is considered to be critical. The number is displayed in red. ■ A threat value that is between 30–69 is considered to be medium-risk. The number is displayed in yellow. ■ A threat value that is between 1–29 is considered to be benign. The number is displayed in blue.
HOST IP	The IP address of the host. Click the IP address link to display the **Host profile** page for the host.
THREATS	Displays the name of the top detected security risk and the number of threats detected on the host. If the name has a icon, click it and a pop-up window displays the description of the threat.

Column Name	Description
THREAT ACTIVITY	Timestamps from when the first event and last event comprising this incident was seen.
CAMPAIGNS	The icon indicates the number of campaigns to which the host belongs.

Host Summary Sidebar

Click anywhere in an entry row for a host in the Host list and the **Host Summary** sidebar appears on the right side of the **Hosts** page.

The following describes what you see in the Host Summary sidebar.

Top Section

The following items are displayed at the top of the panel.

- To close the sidebar, click the ✖ icon.
- The impact value and IP address of the selected host is displayed.
- Point to the impact value and Threat status is displayed.
- To go to the **Host profile** page, click **View Profile**.
- The number of campaigns, threats, applications, and services are displayed.

Details Section

The following details about the host are displayed:

- The Host Name section lists all known host names for the host.
- The Host Label section lists any labels assigned to the host. You can edit the label.

Active Campaigns

The Active campaigns section lists the campaigns associated with this host during the current time frame, if any. Each entry is a summary of a campaign and includes the following information.

- Impact of the campaign.
- The Campaign ID, which is a link to the **Campaign Details** page. See Understanding the Campaign Details Page.
- The number of hosts that are part of the campaign.

Threats

The Threats section lists the threat incidents associated with the selected host during the current time frame. Each entry is a summary of a threat:

- Impact value of the threat.
- The name of the threat. Hovering over the name displays a pop-up with further information about the threat.
- The threat activity time range.

Click the **View threats** link to view the details on the **Host profile > Threats** tab.

Host Profile Page

The **Host profile** page provides an overview and details about the host you selected from the hosts list in the NSX Network Detection and Response **Hosts** page.

The **Host profile** page consists of the following tabs.

Tab Name	Description
Overview	Provides a summary of the host and is the default view.
Threats	Displays the detected incidents, with their associated evidence, network interactions, and IOCs.
Events	Displays detection and info events information.
File downloads	Lists the files that have been downloaded.

There are controls and buttons along the top of the **Host profile** page that are common to all of the tabs.

- Click < to return to the **Hosts page** listing.

  Beside the navigation element is the threat level indicator for the host followed by its IP address. If the host is within the home network, the icon is displayed.

- To launch the **Manage alert** sidebar, click **Host Actions** in the upper-right side of the UI, and select **Manage alert** from the drop-down menu. The **Mange Alert - Filters** sidebar appears on the right side.

  Use the **Manage Alert** sidebar to suppress or demote alerts thrown by harmless events from the host, such as the system Test or Blocking events, or to assign custom impact values to events. See Working with the Manage Alert Sidebar for details.

Host Profile: Overview Tab

The **Overview** tab in the **Host Profile** page in the NSX Network Detection and Response user interface provides a summary about the selected host.

Host Summary

The **Host summary** section contains the **Threats** widget, which provides a quick overview about the threats detected on the host.

Related Campaigns

The Related Campaigns section lists campaigns that affect the selected host. Click the Campaign ID link and the Campaign summary sidebar displays an overview of the campaign.

Host Identity

The Host Identity section contains the following details.

- Host IP - The IP address of the host.
- Host Name - The detected name of the host.
- Host Label - The label for the host. To edit the label, click the icon.

Host Configuration

The Host Configuration section contains the following properties.

- On Home Network - To add the host, click the toggle to **YES**. Otherwise click the toggle to **NO**.
- Silenced - To add the host, click the toggle to **YES**. Otherwise click the toggle to **NO**.

Host Properties

The Host Properties section contains the following details.

- First Seen - Timestamp indicating when the host was first seen.
- Last Seen - Timestamp indicating when the host was last seen.

Host Profile: Threats Tab

Threats detected by NSX Network Detection and Response are represented by threat cards on the **Threats** tab of the **Host Profile** page.

A threat card displays the calculated threat score, the threat name and class, the detection outcome (if available), the threat status, and other actions. If available, the campaign to which this threat is connected is displayed. Expand the card to see its related evidence.

Use the **Sort by** drop-down menu to sort the threat cards. You can select from **Most recent**, **Earliest**, **Highest impact** (the default), and **Lowest impact**.

The **Search threats** text box provides fast, as-you-enter search. It filters the rows in the list, displaying only those rows that have text, in any field, that matches the query string that you provided.

Toggle the **Show closed threats** button to filter the displayed threat cards by threat status. The default is to show all threats.

Managing the Threat Cards

The Threat cards show all the threats associated with the selected host and their corresponding threat levels. Each card displays the calculated threat impact, the threat name, the threat class, and if available, the detection outcome. It also shows the status of the threat: Open or Closed.

Click **Next steps** and select an action from the drop-down menu.

- Select **Close** to close the threat. Select **Open** to reopen a closed threat.
- Select **Manage Alert** to create an alert management rule from the threat.

The **Evidence Summary** section contains an overview of the evidence and other data detected for the threat. Click the › or almost anywhere else in the card to expand the evidence details.

If campaign data connected to this threat is available, **Campaign** with a link to the **Campaign summary sidebar** is displayed.

Evidence details

The **Evidence** column displays the file downloads, signatures, and other categories of evidence type, along with a timestamp of when the evidence was seen. When you click the evidence type link, the corresponding **Evidence Summary** sidebar for that type is displayed on the right side of the page. The **Evidence Summary** sidebar is available for the following evidence types.

- Anomaly
- File download
- Signature

The **Network interactions & network IOCs** column displays the IP address or domain name of external hosts. Clicking the link expands the **Network Interaction** sidebar.

The **Supporting data** column provides a link to the detection events, as well as a link to the threat details.

Detection outcomes

Threat detection event outcomes have the following possible values, listed in order of severity.

Detection Outcome	Description
Succeeded	The threat was verified to have reached its goal. This could be its check-in attempt to the C&C server completed and data was received from the malicious endpoint.
Failed	The threat failed to reach its goal. This could be caused by the C&C server being offline, the attacker made coding errors, and so on.
Blocked	The threat was blocked by the NSX Network Detection and Response application or by a third-party application.

If the event outcome is unknown, this field is not displayed.

Network Interaction Sidebar

You expand the **Network interaction** sidebar by clicking the IP address or domain name link for a specific host in the **Network interactions & network IOCs** column of the **Threats** tab.

The impact and IP address of the selected host are displayed at the top of the sidebar.

WHOIS summary

The **WHOIS summary** section displays key fields from the WHOIS record for the selected IP address or domain name. Click the icon to access the **WHOIS pop-up** window for more details about the IP address or domain. See WHOIS Pop-Up Window for details.

Open in

The **Open in...** section contains links to third-party providers such as DomainTools, VirusTotal, Google, and others. If there are more providers than fit in the view, you can click **Expand for more** to see them.

Anomaly Evidence Summary Sidebar

The **Evidence Summary** sidebar for an evidence type of Anomaly displays when you click an Anomaly evidence link in the Evidence column of the **Threats** tab.

Click **Ref Event** to access the **Event profile** page and the full details of the associated event.

A brief description of the evidence is provided.

Threat details

The following details about the threat are provided.

- Threat – Name of the detected security risk.
- Threat class – Name of the detected security risk class.
- First seen ↔ Last seen – A graph with the timestamp from when the evidence was first and last seen. The duration is displayed below the graph.

Detector summary

A summary of the detector is displayed. For more details, click the **More details** › link to view the **Detector Pop-Up** window. See Detector Documentation Pop-Up Window for details.

- Detector name – The name of the detector.
- Goal – Short description of the goal of the detector.
- ATT&CK categorization – If applicable, a link to the MITRE ATT&CK technique is provided. Otherwise, `N/A` is displayed.

Anomaly details

Details about the anomaly are provided.

Detail	Description
Description	A brief description of the anomaly detailing how it deviates from baseline behavior or why it should be considered suspicious.
State type	The type of anomaly. For example, Outlier.
Anomaly	The anomalous item seen on the host. For example, access to an unusual port.
Baseline items	The items that are typically seen on this host.
Profile created at	Timestamp for the creation of the baseline.
Profile updated at	Timestamp for when the anomaly was detected.
Outlier diagram	The diagram illustrates the normal data upload/download for the host for comparison to the data transfer that was flagged as anomalous. The following data might be displayed, depending on the detector ■ The upload/download size that caused the anomaly alert to be triggered. ■ The maximum upload/download size before the anomaly alert was triggered. ■ The average upload/download size for the host.

File Download Evidence Summary Sidebar

The **Evidence Summary** sidebar for an evidence type of File Download is displayed when you click a File Download evidence link in the Evidence column of the **Threats** tab.

Click **Ref Event** › to access the **Event profile** page and the full details of the associated event.

A brief description of the evidence is provided.

File details

The following details are provided about the file.

- File type – The high-level type of the downloaded file. See Unique Tab for the list of file types.
- Confidence – Indicates the probability that the downloaded file is indeed malicious. As the system uses advanced heuristics to detect unknown threats, in some cases, the detected threat might have a lower confidence value if the volume of information available for that specific threat is limited.
- SHA1 – The SHA1 has of the file.

Malware identification

A summary of the detected malware is displayed. For more details, click the **Analyst report** link to view the Analysis report. See Using the Analysis Report for more details.

- Antivirus class – A label defining the antivirus class of the downloaded file.
- Antivirus family – A label defining the antivirus family of the downloaded file.
- Malware – A label defining the malware type of the downloaded file. If the label has the icon, click the icon to see the description in a pop-up window.
- Behavior overview – The detected behaviors of the downloaded file. If there is a lot of data, a partial list is displayed by default. Click **Expand for more** to view more. Toggle it closed again by clicking **Collapse for less**.

Open in ...

To open the downloaded file in a specific service, click one of the icons for the providers. By default, this displays a partial list of providers.

Download details

The details of the downloaded file are displayed. For more details, click the **Analyst report** link to view the Analysis report. See Using the Analysis Report for more details.

Info	Description
File name	The resource path to the downloaded file.
URL	The full URL to the downloaded file.
First seen	The timestamp from when the downloaded file was first seen. If there have been multiple instances of this file, this will be a range of timestamps.
Downloaded from	The IP address of the source server.
Protocol	The protocol that was used to transfer the downloaded file from the source server.
User agent	If available, the user agent string seen for the download request.

Signature Evidence Summary Sidebar

The **Evidence Summary** sidebar for an evidence type of Signature is displayed when you click a Signature evidence link in the Evidence column of the **Threats** tab.

Click **Ref Event** › to access the **Event profile** page and the full details of the associated event.

A brief description of the evidence is provided.

Threat details

The following details are provided about the threat.

Detail	Description
Threat	Name of the detected security risk.
Threat class	Name of the detected security risk class.
Activity	If available, displays the detected current activity of the threat.
Confidence	Indicates the probability that the detected threat is malicious. For events that show analysis results, for example, a file download, a score is displayed.
First seen ⟷ Last seen	A graph with the timestamp from when the evidence was first and last seen. The duration is displayed below the graph.

Traffic details

The **Reference event traffic** widget provides an overview of the traffic observed between the hosts involved in the referenced event. At least one host involved in the event is a monitored host. The communicating host may be a monitored host or an external system.

The arrow indicates the traffic direction between the hosts.

For each host, the IP address is displayed. If the host is local, the address is a link that you can click to view the Host profile page. A geo-located flag, , or icon might be displayed. More than one can be displayed. If available, a host name is displayed. Any host tags applied to the host are displayed. If available, click the icon to view host details in the **WHOIS** pop-up window. See WHOIS Pop-Up Window for details.

Detector summary

A summary of the detector is displayed. For more details, click the **More details** › link to view the **Detector Pop-Up** window. See Detector Documentation Pop-Up Window for details.

- Detector name – The name of the detector.
- Goal – Short description of the goal of the detector.
- IDS Rule – Click the **View rule (if available)** link to display the **Detector Pop-Up** window. See Detector Documentation Pop-Up Window for details. It can contain an IDS rule.

Host Profile: Events Tab

The **Events** tab in the **Host Profile** page displays detection and events information.

Detection events

The Detection events list shows the events that the NSX Network Detection and Response application found associated with the selected host. These events make up some of the incidents also listed for the host.

Customize the number of rows that are displayed. The default is 30 entries. Use the ‹ and › icons to navigate through the multiple pages.

The columns to be displayed in the list can be customized by clicking the ☰ icon.

Each row displays a summary of an event. Click anywhere on an entry row to access the **Event Summary** sidebar.

The Detection events list contains the following columns.

Column Name	Description
Timestamp	Indicates the start time of the event. The time is shown in the currently selected time zone. The list is sorted by timestamp, by default in decreasing order (latest event at the top). You can use the icons to sort the list in increasing order (oldest event at the top) or toggle back to the default.
Host	The host in the monitored network that is involved in this event. This column will display the IP address, host name, or label of the host, depending on your current Display settings.
Other IP	IP address and port of the host that is related to this event. For example, 203.0.113.115:80 indicates that the IP address 203.0.113.115 was contacted on port 80. The system attempts to geo-locate the IP address. If it succeeds, a small flag icon indicates the country that possibly hosts that IP address. A Local Network icon is used for local hosts.
Other Host	The host name or IP address of the malicious/suspicious entry.
Threat	Name of the detected threat class.
Threat Class	Name of the detected threat class.
Impact	The impact value indicates the critical level of the detected threat and ranges from 1 to 100: ■ Threats that are 70 or above are considered to be critical. ■ Threats that are between 30-69 are considered to be medium-risk. ■ Threats that are between 1-29 are considered to be benign. If the ⊘ icon appears, it indicates the artifact has been blocked. Click the ⬍ icon to sort the list by impact.

Info Detection Events

The Info Detection Events list shows `INFO` events associated with the selected host. This list contains the same columns as the Detection events list.

Host Profile: File Downloads Tab

The **File downloads** tab in the **Host Profile** page of the NSX Network Detection and Response UI shows the malicious files downloaded by the host with details about their contents and corresponding threat levels

The **Quick search** text box above the list provides fast, as-you-enter search. It filters the rows in the list, displaying only those rows that have text, in any field, that matches the query string.

The columns to be displayed in the list can be customized by clicking the ☰ icon.

Each row is a summary of a downloaded file. Click the ⊕ icon (or anywhere on an entry row) to view details of the downloaded file.

The list is sorted by score and includes the following columns.

Column Name	Description
Timestamp	The timestamp of the detection of the file download
Host	The host that downloaded the file.
Sensor	The sensor that detected the file download.
Contacted IP	The IP address of the contacted host.
Location	For a download, this is the URL of the file in the supported format. For example, `\\127.0.0.2\samba_share\1128dedb.exe` for an SMB download or `http://www.example.com/download/example.zip` for an HTTP download. For an upload, "Upload" is displayed.
Filename	The name of the file downloaded.
MD5	The MD5 hash of the downloaded file.
Type	The high-level file type of the downloaded file. See Unique Tab for the list of currently supported types.
AV Class	A label defining the antivirus class of the downloaded file. If the label has a  , you can click that icon for a description in a pop-up window.
Malware	A label defining the malware type of the downloaded file. If the label has a  , you can click that icon for a description in a pop-up window.
Score	The score assigned to the downloaded file by the analysis indicates the critical level of the detected threat and ranges from 0–100: ■ Threats that are 70 or higher are considered to be critical. ■ Threats that are between 30–69 are considered to be medium-risk. ■ Threats that are between 1–29 are considered to be benign. For details about maliciousness core and risk estimate, see Analysis Report: Overview Tab. If the  icon appears, it indicates the artifact has been blocked. The list is sorted by decreasing order (most critical threats at the top). Click  to sort the list in increasing order (least critical threats at the top), then click  to toggle back to the default.

Working with the Events Page

The **Events** page provides information on individual events that the NSX Network Detection and Response application detected in your NSX network.

The page consists of several widgets that can be managed using the information in Getting Familiar with the NSX Network Detection and Response User Interface.

The **Network** tab on the **Events** page consists of widgets that allow you to inspect, manage, and prioritize the network detection events reported by the NSX Network Detection and Response application.

Global Event Map

The **Global event map** widget provides a visual overview of geolocations of the aggregated events.

It marks the approximate location of the other hosts involved in the event detected by the NSX Network Detection and Response application. The marker color represents the event impact. The marker size represents the number of impacted hosts.

Events with no specific location are excluded from this map.

To learn more about the threats and hosts represented at that particular location, click a marker on the map.

In the **Location Details** pop-up window that displays, you can view the approximate location, the threats, and the destination hosts for the selected event. Click the ▼ icon next to each entry to apply filters to the list displayed in the **Events** page.

Detected Threats in the Events Page

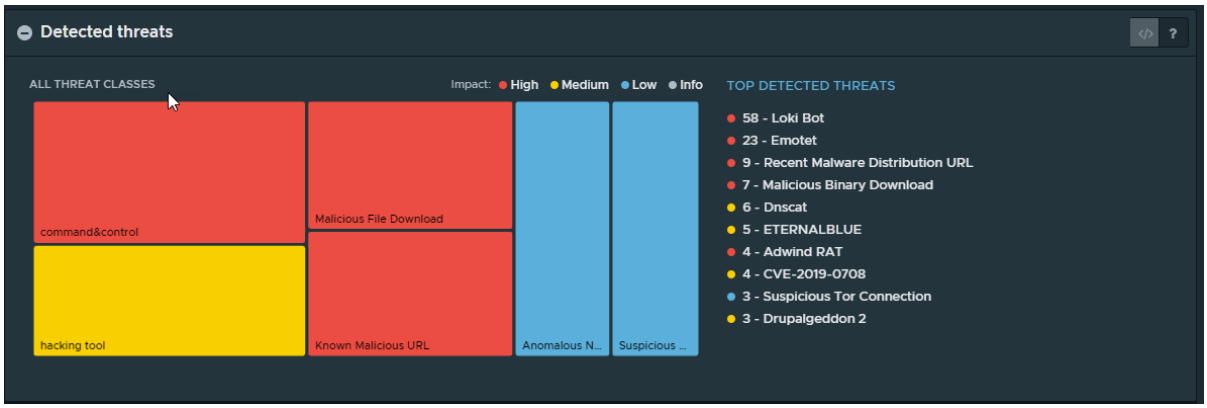
The **Detected threats** widget on the **Events** page provides a visualization of all types of threat classes and threats that the NSX Network Detection and Response application discovered in your network.

By clicking the rectangle for a specific threat class, you can further examine the threats it contains within the same visualization. When you select a specific threat, the system displays details about that particular threat and its activity in your network.

Note Your selections, as you navigate to the individual threats, trim the **Detection Events** list. Conversely, when you use the filters to narrow the displayed list of events, the threats presented in the **Detected threats** widget are also filtered.

Threat Class

The initial view shows the threat classes that have been detected on your network, similar to the following image.



The rectangles represent the threat classes that have been detected on your network. The size of each rectangle is scaled based on the number of events for each detected threat class. The colors of the blocks indicate the severity of the threat.

The list on the right side of the widget shows the list of top detected threats. When you point to an item in the list, a pop-up window gives further information about the threat, its class, and the number of events and affected hosts.

When you point to a specific rectangle for a threat class, a pop-up window appears. It shows the threat class, the number of unique threats, and a breakdown of the number of events and participating hosts. Clicking the pop-up window or the rectangle allows you to drill down into the unique threats that make up the selected threat class.

Unique Threats

The subsequent view shows the threats that make up the selected threat class. The rectangles are scaled based on the number of events for each detected threat and the colors indicate the severity of the threat.

A pop-up window is displayed when you hover over a specific threat. It shows the threat and a breakdown of the number of events and participating hosts. When you click the pop-up window or the rectangle to select the threat, **Threat Details** is displayed on the right side of the widget.

Threat Details

The Threat details section lists the following information:

- THREAT: The name of the threat.
- CLASS: The name of the threat class.
- MAX IMPACT: The maximum impact of events detected for the threat.
- EVENTS: The number of detected events.
- HOSTS: The number of targeted hosts. To view the Hosts list, click the number link. See Hosts List for more details.
- FIRST SEEN/LAST SEEN: A bar graph that shows the timestamps seen for the threat. The Duration is displayed underneath.

Use Filters on the Events Page

The NSX Network Detection and Response application provides a filtering mechanism that allows you to focus on specific events information that is of interest to you. The use of filters is optional.

Procedure

1. From the **Events** page, click  to expand the **Filters** widget.
2. Click anywhere in the **Filter on** text box and select an item from the drop-down menu.

   You can select from the following available filters. To further narrow the focus of the displayed information, you can combine multiple filters.

Filter Name	Description
Event outcome	Select **All** or **Info** from the drop-down menu. The default is to display events that are determined to be related to a threat. Selecting **Info** includes only those events that themselves are informational. By tracking these events, you can gain further insight into the activity in your network.
Home network	Restrict displayed events by the Home network setting using the drop-down menu. Select **Home network only** for events within your defined home network. Select **Unidentified networks only** for events from unknown hosts.
Host IP	Restrict displayed events to a specific source IP address, IP address range, or CIDR block. Enter a valid value in the **Host IP** text box.
Host name	Restrict displayed events to a specific source Host name. The full host name or label needs to be provided.
Incident ID	Display events that belong to the specified Incident. An Incident ID is a numeric entry, for example, `73142`. A valid incident ID must be provided.
Minimum impact	Display events that scored the minimum impact level. The range is 1–100.
Other host	Restrict the displayed events to a specific host name.
Other host IP	Restrict the displayed events to a specific host IP address. The IP address can be entered as one or more IP addresses, CIDR blocks (such as `192.168.0.0/24`) or IP address ranges (such as `1.1.1.5-1.1.1.9`).
Port	Display events using a specific TCP/UDP port. To further filter the displayed events, you can combine this with the **Transport** filter.
Priority	Restrict displayed events by the Priority status. Select **Infections**, **Watchlist**, or **Nuisances** from the drop-down menu. See Infections Over Time for details.
Threat	Restrict displayed incidents by a specific Threat. Select a threat from the drop-down menu. The menu is prepopulated with a list of cataloged threats. Use the search function at the top of the menu to quickly find a threat name.
Threat class	Restrict display to a specific class of events. Select the threat class from the drop-down menu. The menu is prepopulated with a catalog of classes.
Transport	Display events using a specific transport layer protocol. Select **TCP** or **UDP** from the drop-down menu.

3 To apply the selected filters, click **Apply**.

The system applies the selected filters and updates the Events list.

4 (Optional) To delete an individual filter, click the **REMOVE–** button next to its entry. To delete all the selected filters, click the **X** icon located on the right side of the **Filter** widget.

The **Filters** widget collapses when you delete all the selected filters.

Detection Events

The **Detection Events** widget provides an overview of the individual events that the NSX Network Detection and Response application detected.

An event represents a security-relevant activity that has occurred in the monitored network. An event may involve multiple data flows (for example, TCP connections), but it represents a single type of activity occurring over a short period of time (at most one hour).

If the selected time range includes today (the default), the widget updates its list of events every 5 minutes. New events are highlighted in green; the color fades away after a few seconds.

The **Quick search** field above the list provides fast, as-you-enter search. It filters the rows in the list, displaying only those rows that have text, in any field, that matches the query string.

Manually refresh the events list by clicking the **Update Now** button.

Customize the number of rows to be displayed. By default, 30 entries are shown. Up to 1000 events can be displayed, however, there may be a noticeable delay for the system to retrieve a large number of events. Use the ‹ and › icons to navigate through multiple pages.

Each row displays a summary of an event. Click anywhere on an entry row to access the **Event Summary** sidebar.

The list of events contains the following columns.

Column Name	Description
Timestamp	Indicates the start time of the event. The time is shown in the currently selected time zone. The list is sorted by timestamp, by default in decreasing order (latest event at the top). You can use the icons to sort the list in increasing order (oldest event at the top) or toggle back to the default. Click the ⬍ icon to sort the list by timestamp.
Host	The host in the monitored network that is involved in this event. This column will display the IP address, host name, or label of the host, depending on your current Display settings. Click the **Edit** icon next to the host to open the Label/Silence host pop-up.
Other IP	IP address and port of the host that is related to this event. For example, 203.0.113.115:80 indicates that the IP address 203.0.113.115 was contacted on port 80. The system attempts to geo-locate the IP address. If it succeeds, a small flag icon indicates the country that possibly hosts that IP address. A Local Network icon is used for local hosts.
Other Host	The host name or IP address of the malicious/suspicious entry.
Threat	Name of the detected threat or security risk.
Threat Class	Name of the detected threat class.
Impact	The impact value indicates the critical level of the detected threat and ranges from 1 to 100: ■ Threats that are 70 or above are considered to be critical. ■ Threats that are between 30-69 are considered to be medium-risk. ■ Threats that are between 1-29 are considered to be benign. If the ⊘ icon appears, it indicates the artifact has been blocked. Click the ⬍ icon to sort the list by impact.

Event Summary Sidebar

You access the **Event Summary** sidebar when you click an entry row in the **Detection Events** widget of the NSX Network Detection and Response**Events** page.

The following section describes what you see on this sidebar. After the top section, subsequent sections display supporting data. Some sections are displayed only if relevant data is available.

Top section

The top of the sidebar includes the following:

- To close the sidebar, click the ✖ icon.
- To view the event in the **Event profile** page, click **Details** ›. See Event Profile Page for more information.
- If available, a brief description of the event is provided. It includes an explanation as to why the system flagged this event, identifies the threat or malware associated with this event, and briefly describes the detected activity.

Threat Details

This section includes the following information.

Threat Detail Name	Description
Threat	Name of the detected security risk.
Threat Class	Name of the detected security risk class.
Event Detector	The name of the event detector. Click the link to view the **Detector** pop-up window. See Detector Documentation Pop-Up Window for details. If there is no detector for the event, this section is not shown.
Impact	The impact value indicates the critical level of the detected threat and ranges from 1-100 ■ Threats that are 70 or above are considered to be critical. ■ Threats that are between 30-69 are considered to be medium-risk. ■ Threats that are between 1-29 are considered to be benign.
Action	A list of actions taken by the sensor (for example, any blocking activities, whether the event is logged, whether traffic was captured, or a malware download was extracted).
Outcome	The outcome of the event. In most cases, this is Detection. For Info events and events that were promoted from Info status, an additional label provides the reason for its status/status change. A pop-up window is displayed when you hover over the label, providing additional details about the reason.
First Seen Last Seen	A graph with the timestamp from when the evidence was first and last seen. The Duration information is displayed below the graph.

Event traffic

The **Event traffic** widget provides an overview of the traffic observed between the hosts involved in the event. At least one host involved in the event is a monitored host. The communicating host can be a monitored host or an external system. A link to view the Captured traffic is displayed, if the data is available.

The arrow indicates the traffic direction between the hosts.

For each host, the IP address is displayed. If the host is local, the address is a link that you can click to view the **Host profile** page. A geo-located flag, , or might be displayed. More than one might be displayed. If available, a host name is displayed. If available from DHCP traffic monitoring, the MAC address of the host is displayed. Any host tags applied to the host are displayed. If available, click to view host details in the **WHOIS** pop-up window.

Event evidence

The Event evidence section lists various actions observed while analyzing the event. For more details, click the **Event details** link to view the Event evidence.

Actions include Signature, Reputation, Unusual behavior, File download, URL path match, Verification, Anomaly, and so on. If provided, click the link to view the corresponding **Detector** pop-up window A Confidence value is displayed for each action.

Malware identification

If the NSX Malware Prevention application is activated, a summary of the detected malware is displayed. For more details, click the **Analyst report** › link to view the Analysis report. See Using the Analysis Report for more information.

Detail Name	Description
Antivirus Class	A label defining the antivirus class of the downloaded file.
Antivirus Family	A label defining the antivirus family of the downloaded file.
Malware	A label defining the malware type of the downloaded file. If the label has a icon, you can click it for a pop-up description.
Behavior Overview	The detected behaviors of the downloaded file. If there is a lot of data, a partial list is displayed by default. Click **Expand for more** to view more. Toggle it closed again by clicking **Collapse for less** .

Event URLs

The Event URLs section displays all the URLs detected in the event. This section appears only if the event is associated to a URL

Event metadata

The Event metadata section displays the following data.

Data Name	Description
Related Incident	Click to view the related incident, if one is available.
Connections	The number of connections included in the event.
Related Campaign	Click to view the related campaign, if one is available.

WHOIS Pop-Up Window

The **WHOIS** pop-up window displays registration information and other details about the IP address or hostname of the host you are examining.

It has the following two tabs.

Summary

The **SUMMARY** tab displays the following information about the IP address or hostname.

- Date information – The date the domain was registered, the date that the domain record was updated, and if available, the expiration date of the domain.
- Organization – The name of the organization, the organization email addresses, the organization country (country code), the organization phone numbers, the registrar name, and the contact list.
- Network – The network name, IP address range, AS list,authoritative name servers, and parent networks.

Raw Record

The **RAW RECORD** tab displays WHOIS data in its raw form.

Information unavailable

If the **WHOIS** pop-up window displays a warning that information for the given IP address or hostname is unavailable, you can try using a 3rd party. To look-up the host, click **View in external tool** in the bottom-right of the pop-up window.

Note The button to the 3rd party provider is always available.

Detector Documentation Pop-Up Window

The **Detector documentation** pop-up window provides detailed information about the NSX Network Detection and Response detector that provided the event evidence. The intent is to assist you in determining the confidence you can place in this detector.

The documentation displays at least some of the following details.

Detail Name	Description
Goal	Short description of the goal of the detector.
ATT&CK categorization	If applicable, a link to the MITRE ATT&CK technique is provided.
Detector abstract	A detailed technical description of the detector and its operation.
IDS rule	A high-level representation of the detection logic used by an NSX Network Detection and Response network signature. The rule syntax is loosely related to the Suricata signature language defined in the Suricata Rules documentation. A rule consists of one or more clause sets, typically a single clause, each containing key/value pairs. If there is more than one clause in a rule, each clause is numbered. The first clause is prefaced "IF:" and each subsequent clause is prefaced with "AND THEN IF:". The different clause sets are evaluated sequentially on data belonging to the same flow. Point to any key/value pair to view a relevant help pop-up window.
False positives	A description of the possibility of the detector to generate false positives.
False negatives	The assumptions that might result in the detector causing false negatives.

Event Profile Page

The **Event profile** page is accessed from the **Details** › button at the top of the **Event Summary** sidebar.

There are a number of controls and buttons along the top of the view:

- Click **Similar Events** to view a drop-down list of similar features. Click the ☑ icon beside each to select **Destination**, **Destination port**, **Source IP**, **Transport protocol**, **Threat class**, and **Threat type**. Then click **View Events** to view the selected events in a new tab.
- Click **Manage Alert** to launch the **Manage alert** sidebar. Use this feature to suppress or demote harmless events, such as the system Test or Blocking events, or to apply custom scores to specific events. See Working with the Manage Alert Sidebar for details.
- Click the ⊖ icon to collapse all fields or the ⊕ icon to expand all fields.

Event Overview

The top section provides a visual overview of the threat or malware that the NSX Network Detection and Response application detected, and displays the threat class and the threat impact score.

Event Summary

The **Event summary** section provides an explanation as to why the NSX Network Detection and Response application flagged this event, identifies the threat or malware associated with this event, briefly describes the detected activity, and displays supporting data.

If available from the NSX Advanced Threat Prevention cloud service, a detailed explanation of the event and why it is considered malicious is displayed at the top of the **Event summary** section.

Server Block

The Server block displays the following data.

Data	Description
Host name	If available, the FQDN of the server.
IP address	The IP address of the server. A geo-located flag might be displayed. If the icon is present, click the link to view more details in the **Host profile** page. If available, click the icon to see the reputation tags of the client. If available, click the icon to view registration information and other data about the host in the **WHOIS** pop-up window.
MAC address	If available, the MAC address of the server. This address is obtained from monitoring DHCP traffic and is one of the data points the system uses to generate a unique HostID entry that it maps to a specific host in the network, regardless of its IP address.

Client Block

The Client block displays the following data.

Data	Description
Host name	If available, the FQDN of the client.
IP address	The IP address of the client. A geo-located flag might be displayed. If available, click the address or the icon to view the **Host profile** page . If available, click the icon to see reputation tags of the client. If available, click the icon to view registration information and other data about the host in the **WHOIS** pop-up window.
MAC address	If available, the MAC address of the client. This address is obtained from monitoring DHCP traffic and is one of the data points the system uses to generate a unique HostID entry that it maps to a specific host in the network, regardless of its IP address.

Event Metadata

The Event metadata section displays the following data.

Data	Description
Verification outcome	Indicates the event outcome. The following are the possible values. ■ Blocked: The threat was blocked by the NSX Network Detection and Response application or by a third-party application. ■ Failed: The threat failed to reach its goal. This could be caused by the C&C server being offline, the attacker made coding errors, etc. ■ Succeeded: The threat was verified to have reached its goal. This could be its check-in attempt to the C&C server completed and data was received from the malicious endpoint. If the event outcome is unknown, this field is not displayed.
Verifier name	The name of the event verifier. Click the link to access the **Verifier Documentation** pop-up window.
Verifier message	A message from the verifier which provides further information about the outcome, for example, which third party application blocked the threat.
Sensor	The sensor that detected the event.
Connections	The number of connections included in the event.
Action	A list of actions taken by the sensor (for example, any blocking activities, whether the event is logged, whether traffic was captured, or a malware download was extracted).
Users logged in	A list of the users detected in the logged records.
Outcome	The outcome of the event. In most cases, the outcome is `DETECTION`. For INFO events and events that were promoted from INFO status, an additional label provides the reason for its status/status change. A pop-up is displayed when you hover over the label, providing additional details about the reason.
Related incident	A permalink to a correlated incident. Clicking the link opens the **Incident profile** page in a new browser tab. This event might be one of a number of closely related events that have been automatically correlated into an incident.
Event ID	View the event in the **Network event details** page. The link opens in a new browser tab.
Start time	A timestamp for the beginning of the event.
End time	A timestamp for the end of the event.

Captured Malware

The Captured malware section provides information from the dynamic analysis that was performed on the malicious software instance that is related to the event.

You can access detailed in-depth technical information on what the malware does, how it operates, and what kind of a risk it poses. For more information on the displayed information, see Using the Analysis Report.

Note If no malicious software was detected for the event, this section will not appear.

Event Evidence

The Event evidence section provides details of the actions observed while analyzing the event.

Actions can include malicious file download, network traffic matching the network signature for known threats, performing a domain name resolution of a blocked malware domain, a known bad URL path, and so on.

If available, click the Detector link to view the Detector Documentation Pop-Up Window pop-up window. Also see the About Evidence for more details.

Host Reputation

The Host reputation section provides information about known malicious hosts or URL reputation entries seen in the event.

Note If the host has no known history, this section will not appear.

Anomaly Data

This section displays the netflow or passive DNS records that caused the anomaly event to be raised.

It will be titled **DNS anomaly data** or **Netflow anomaly data**, depending upon the anomaly seen.

Additional information may be provided, such as the IP addresses or ports that have been classified as anomalous. If a large number of items are involved, you can click the +# to expose all the items.

Note If no anomalies were seen for the event, this section will not appear.

Threat Description

The Threat description section provides a detailed description of the threat associated with the event.

Mitigation

The Mitigation section provides detailed instructions for the removal of any malicious software and other recommended processes to clean up after the event.

Note If there is no known mitigation process for the event, this section will not appear.

Managing the Incidents Page

The **Incidents** page displays the incidents and their different threat ratings. You can use the widgets in the page to inspect, manage, and prioritize the incidents reported by the NSX Network Detection and Response application.

The page consists of several widgets that can be managed using the information in Getting Familiar with the NSX Network Detection and Response User Interface.

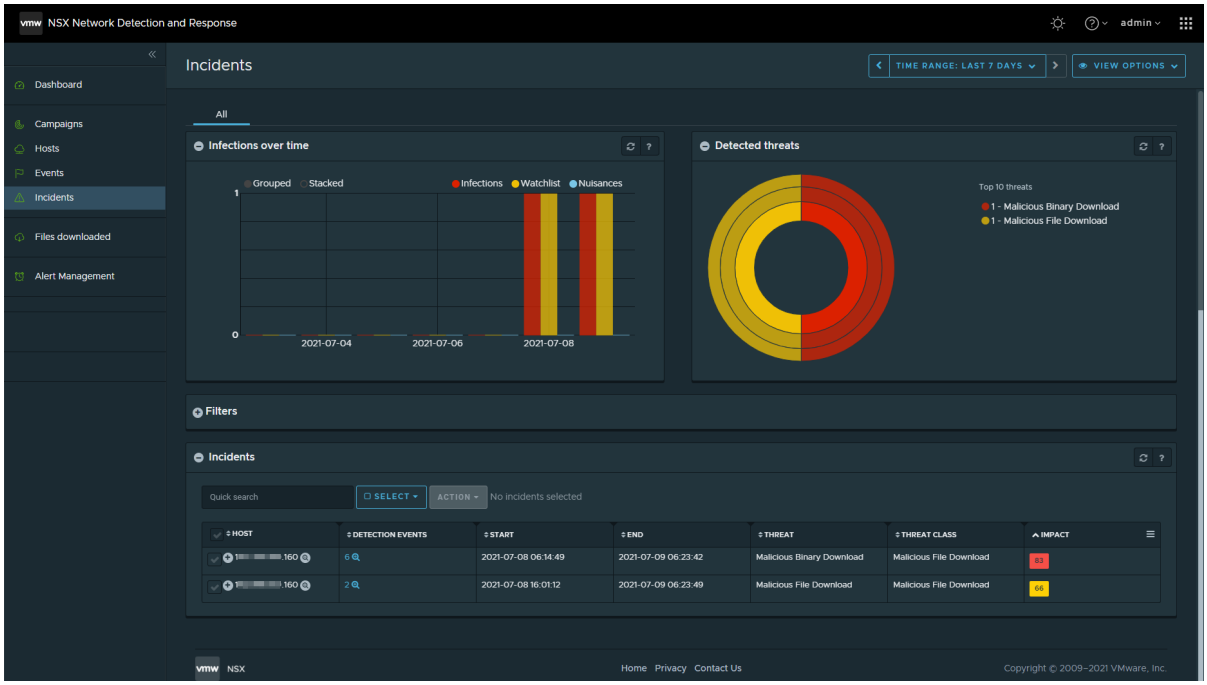
In the NSX Network Detection and Response application, an incident is an aggregation of detection events from a single threat detected on a single workload in the monitored network.

The NSX Network Detection and Response application does not solely report security events. An incident can consist of a single event, or many events that have been automatically correlated and determined to be closely related by the system threat engine. For example, the **Incidents** page can report all outgoing connections to the command and control channel of the malware, all suspicious DNS look ups (for example, requests for automatically generated related malware domains), and in-depth descriptions of each registered security event.

The **Incidents** page allow you to perform the following tasks.

- Efficiently keep track of all incidents that are occurring.
- Quickly see a list of affected hosts.
- Prioritize threats according to their impact and severity levels using different views.
- Gain an in-depth understanding of the events that have been registered for each incident, and access threat and mitigation descriptions.
- Close or open incidents.
- Mark or clear affected hosts as being cleaned.
- Filter reported threats for specific hosts.

The following image is an example of the **Incidents** page, with the **All** tab in display.



Infections Over Time

The **Infections over time** widget provides a graphical overview of the different kinds of incidents detected in the network. The x-axis depicts the time and the y-axis the number of hosts affected by incidents of a given type.

There are three different types of incidents.

Incident Type	Description
Infections	These are incidents that have been determined to be critical. These incidents have been given an impact score of 70 or above and are displayed in red
Watchlist	These are incidents that have been determined to be of medium risk. Such incidents, while indicating a potential risk, may not need immediate attention; they are kept under close watch in case new evidence appears that modifies their status. These incidents have been given an impact score of between 30 and 69 and are displayed in orange.
Nuisances	These are incidents that are considered low or no risk. This typically corresponds to potentially unwanted/risky activity that does not necessarily indicate a compromise or infection on the monitored network. These incidents have been given an impact score of lower than 30 and are displayed in blue.

You can display or hide the different incident types by clicking their corresponding names in the legend at the top of the graph.

When you point to a bar on the graph, a pop-up window displays the number of hosts in the network that are affected by the corresponding incidents.

When you click a bar, the time range and incident type is updated accordingly. The dashboard only displays information for that incident type on the selected day.

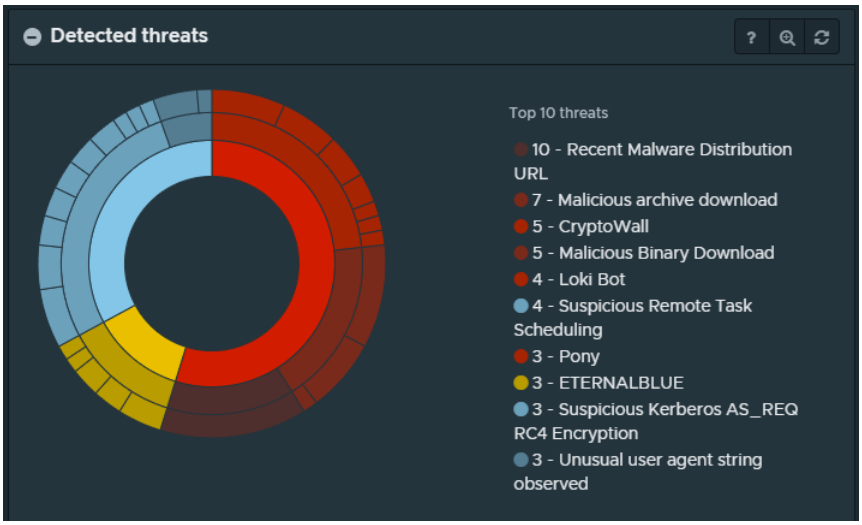
To undo the zoom, reset the time range. Note that this will leave the incident type selected. To reset the dashboard, use the back button in your browser.

The default view shows the incidents in grouped display. Click **Stacked** to view the incidents in a stacked display. Click the **Grouped** to reset to the grouped display.

Detected Threats

The **Detected threats** widget provides a graphical overview of the different kinds of threats that the NSX Network Detection and Response application has detected in the network.

The threat information is displayed in a layered circle, similar to the following image.



The divisions of the circles represent the number of hosts affected by the displayed incident types. Moving toward the outer circles provides a finer granularity and more specific information.

- The innermost ring displays the three different types of incidents.

Incident Type	Description
Infections	These are incidents that the NSX Network Detection and Response application determined to be critical. These incidents have been given an impact score of 70 or higher and are displayed in red.
Watchlist	These are incidents that the NSX Network Detection and Response application determined to be of medium risk. Such incidents, while indicating a potential risk, might not need immediate attention. They are kept under close watch in case new evidence modifies their status. These incidents are assigned an impact score anywhere from 30–69 and they are displayed in yellow.
Nuisances	These are incidents that are considered low or no risk. This typically corresponds to potentially unwanted/risky activity that does not necessarily indicate a compromise or infection on the monitored network. These incidents have been given an impact score of lower than 30 and are displayed in blue.

- The middle ring displays the threat class together with the number of relevant incidents for each type of infection. Threat classes include command&control servers, malicious file downloads, crypto-miners, and many more.

- The outer ring represents the individual threat families detected in the network. Threat families include ransomware, malicious binary files, and so on.

When you point to the graph, the widget displays the threat name and a count of hosts where the NSX Network Detection and Response application observed the threat.

When you click an item in the graph, the view zooms in and displays more details about the selected information type. Clicking the item again zooms the view back.

If you click an incident type in the inner ring, the graph view zooms in and displays the matching incidents in the middle and outer ring. If you click a threat class in the middle ring, the graph view zooms in and displays the matching threat families. If you click the outer ring, the graph view zooms in and displays details about the selected threat.

The legend on the right side of the widget provides a count of the occurrences of the most frequent threats detected. When you point to an item in the legend, a pop-up window gives further information about the threat class, the number of incidents, and the number of affected hosts. Clicking the item zooms the graph view for the selected threat type and provides more contextual information.

Use Filters on the Incidents Page

NSX Network Detection and Response provides a filtering mechanism that allows you to focus on specific incident information that is of interest to you. The use of filters is optional.

Procedure

1 From the **Incidents** page, click ⊕ to expand the **Filters** widget.

2 Click anywhere in the **Filter on** text box and select an item from the drop-down menu.

You can select from the following available filters. To further narrow the focus of the displayed information, you can combine multiple filters.

Filter Name	Description
Campaign UUID	Restrict the displayed entries by the Campaign UUID. This is a 32-character hexadecimal string, for example, `7dabc0fc9b3f478a850e1089a923df3a`. Alternatively, enter the string `null` to select records that do not belong to any campaign.
Home network	Restrict the displayed entries by the Home network setting. Select **Home network only** or **Unidentified networks only** from the drop-down menu.
Host IP	Restrict the displayed entries to a specific source IP address, IP address range, or CIDR block. Enter the value in the text box.
Hosts name	Restrict the displayed entries by the Host name. The full host name or label needs to be provided.
Priority	Restrict the displayed entries by the Priority status. Select **Infections**, **Watchlist**, or **Nuisances** from the drop-down menu.
Read	Restrict the displayed entries by their Read status. Select **Read** or **Unread** from the drop-down menu.

Filter Name	Description
Status	Restrict the displayed entries by their status. Select **Closed** or **Open** from the drop-down menu.
Threat	Restrict the displayed entries by a specific threat. Select a threat from the drop-down menu. The menu is prepopulated with a list of cataloged threats. Use the search function at the top of the menu to quickly find a threat name.
Threat class	Restrict the displayed entries to a specific class of threats. Select the threat class from the drop-down menu. The menu is prepopulated with a catalog of classes, some of which are listed below. Use the search function at the top of the menu to quickly find a class name. ■ **adware**: Malware that displays or downloads advertisements to an infected computer. ■ **click-fraud**: Click-fraud targets pay per click online advertising. ■ **command & control**: An infected machine belongs to a botnet and the machine can be remotely controlled by an attacker. ■ **drive-by**: An attacker attempted to exploit a vulnerability on the machine in order to install additional malware on the target system. ■ **exploit toolkit**: Detection of an exploit toolkit that attempted a drive-by download attack ■ **fake-av**: Fake antivirus software or other kinds of rogue security software designed to trick or mislead your users. ■ **inactive C&C**: The command & control server for this specific botnet is inactive. ■ **VMware blocking test**: The domain block.lastline.com is used to test blocking of network connections and the selected events belong to this class. ■ **VMware test**: The domain test.lastline.com is used to test the functionality of the setup and the selected events belong to this class. ■ **Malicious File Download, Malware Distribution, and malware download**: The IP address or domain hosts malicious executables. ■ **sinkhole**: A sinkhole is operated by a legitimate organization, so it does not pose a threat. However, hosts that try to contact such a host may be infected. ■ **spyware**: Malware that attempts to steal sensitive information. ■ **suspicious-dns**: Suspicious DNS domains are domains that are contacted by malware running on infected machines. Our proprietary techniques were able to proactively identify these domains as malicious. ■ **unknown**: An unknown security risk was detected.

3 To apply the selected filters, click **Apply**.

4 (Optional) To delete an individual filter, click the **Remove –** button next to its entry. To delete all the selected filters, click the ✖ icon located on the right side of the **Filter** widget.

The **Filters** widget collapses when you delete all the selected filters.

Incidents List

An incident represents a security-relevant activity that NSX Network Detection and Response detected has occurred in the monitored network. An incident may consist of a single event, or a number of events that have been automatically correlated, and that have been determined to be closely related. The incidents list shows the registered incidents with their corresponding threat levels.

You can see all reported incidents that have been determined to be critical, those that you should keep an eye on, or those that are considered to be nuisances in your network. Critical incidents must be handled without delay. Failing to deal with critical incidents is highly risky, and increases the probability that other hosts in your network may be compromised as well.

Incidents that you have not examined yet are marked as unread, while those that you have already examined are marked as read. You have the option of selecting incidents and to perform actions on them such as marking them as read or unread. You can also close or open selected incidents.

The **Quick search** text box above the list provides fast, as-you-enter search. It filters the rows in the list, displaying only those rows that have text, in any field, that matches the query string.

Use the **SELECT** drop-down menu for a fine-tuned selection. Its options allow you to select **All visible incidents** or to **Clear selection**. You can also select **Read (current page)** or **Unread (current page)** incidents. You can also click the **Edit** icon in the title row to select all visible messages.

Use the **ACTION** drop-down menu to update the selected incidents: **Mark as read**, **Mark as unread**, **Close**, or **Open**.

Customize the number of rows to be displayed. The default is 20 entries. Use the ‹ and › icons to navigate through multiple pages.

The columns to be displayed in the list can be customized by clicking the additional content icon.

Each row is a summary of an incident. Click the **Plus** icon (or anywhere on an entry row) to access the incident details. To select a message row, click the **Edit** icon.

The list is sorted by Impact and includes the following columns.

Column	Description
Host	The host affected by this incident. This column displays the IP address, host name, or label of the host, depending on the current Display settings pop-up. Click the icon to view the **Host profile** page, showing details about the host. Click the icon to sort the list by host information.
Detection Events	Number of events that comprise this incident. This is a link displaying an event count and the icon. Clicking this link loads the **Events** page, filtered to show only events for this incident. Click the icon to sort the list by events.
Start	Start time of incident. Click the icon to sort the list by start time.
End	End time of incident. Click the icon to sort the list by end time.
Threat	Name of the detected security risk. Click the icon to sort the list by threat.

Column	Description
Threat class	Name of the detected security risk class. Click the **Sort** icon to sort the list by threat class.
Impact	The impact value indicates the critical level of the detected threat and ranges from 1 to 100: ■ Threats that are 70 or above are considered to be critical. ■ Threats that are between 30 -69 are considered to be medium-risk. ■ Threats that are between 1- 29 are considered to be benign. If the stop icon appears, it indicates the artifact has been blocked. The list is sorted by decreasing order of impact (most critical incidents at the top). Click the icon to sort the list in increasing order (least critical incidents at the top), then click the angle down icon to toggle back to the default.

Incident Details

When you click anywhere in an incident row, the Incident Details view is expanded within the incident list.

There are a number of buttons along the top of the incident details:

- Click the button to close the incident.
- Use the **Action** drop-down menu to perform an action on the incident:
  - If the incident is not yet closed, select **Close incident**. Otherwise, select **Open incident**.
  - If the incident is not yet read, select **Mark as read**. Otherwise, select **Mark as unread**.
  - Select **Ignore threat**. The threat details are listed in the menu item. Selecting this item indicates that the presence of this particular threat on the host is not of interest. Therefore, all incidents where this threat is detected on this host are closed automatically.
  - Select **Mark host <host> as cleaned**. The system marks the host that is involved in the incident as cleaned. As a result, all incidents on that host are closed.
- Clicking **View Incident Details** displays the contents of the **Incident Profile** page in a new browser tab.
- Clicking **Manage Alert** launches the **Manage alert** sidebar. Use this feature to suppress or demote harmless events associated with the specified incident, such as the system Test or Blocking related incidents. See Working with the Manage Alert Sidebar for more details.
- Click **Mark as read** to mark the incident. The button toggles to **Mark as unread** which allows you to revert its read status.

Incident summary

The top section provides a visual overview of the detected threat and displays its impact score.

Incident Details

The **Incident Details** widget displays detailed network information about the incident. It includes the following data.

Column	Description
Source IP	The IP address of the incident source. Click the icon to view the **Activity for host** page. Click the icon to view the source in the **Network analysis** page.
Source host	If available, the FQDN of the incident source.
Events	The number of events that make up this incident.
Incident ID	A permalink to the **Incident profile** page. The link opens in a new browser tab/window.
Campaign ID	A permalink to the campaigns page. The link opens in a new browser tab.
Impact	The impact score applied by the system to this incident.
Start time	A timestamp for the beginning of the incident.
End time	A timestamp for the last recorded event of the incident.
Status	Shows if the incident has been closed.

Evidence

The **Evidence** widget when expanded displays the list of events detected by NSX Network Detection and Response.

The columns to be displayed in the list can be customized by clicking the icon.

Each row is a summary of an evidence entry and includes the following columns.

Column	Description
First seen	Timestamp from when this event was first seen.
Last seen	Timestamp from when this event was last seen.
Threat	Name of the detected security risk.
Threat class	Name of the detected security risk class.
Impact	The impact score applied to this incident.
Evidence	The evidence category of this incident. The title of the evidence details block is derived from the category name.
Subject	The artifact, typically a file, that is being analyzed.
Reference	A permalink to the event page. The link opens in a new browser tab.

Evidence Details

Click the icon (or anywhere on an incident entry row) to display the evidence details block.

The title of the evidence details block is derived from the type of evidence. For example, Reputation Evidence.

This section displays more detailed information about the evidence. It includes the following data.

Data	Description
Threat	Name of the detected security risk.
Threat class	Name of the detected security risk class.
Impact	The Impact score applied to this incident.
Detector	If present, displays the NSX Network Detection and Response module that identified the threat. Click the link to view the Detector pop-up window. See Detector Documentation Pop-Up Window.
View network event	A permalink to the event page. The link opens in a new browser tab.
View network event	A permalink to the event page. The link opens in a new browser tab.
First seen	Timestamp from when this event was first seen.
Last seen	Timestamp from when this event was last seen.
Severity	An estimate of how critical the detected threat is. For example, a connection to a command and control server is typically considered high severity as the connection is potentially damaging.
Confidence	Indicates the probability that the detected individual threat is indeed malicious. As the system uses advanced heuristics to detect unknown threats, in some cases, the detected threat may have a lower confidence value if the volume of information available for that specific threat is limited.
Subject	If present, displays the artifact, typically a file, that is being analyzed.

See About Evidence for further details.

Working with the Files Downloaded Page

The **Files downloaded** page provides tabs containing information about the files that were downloaded in your NSX network.

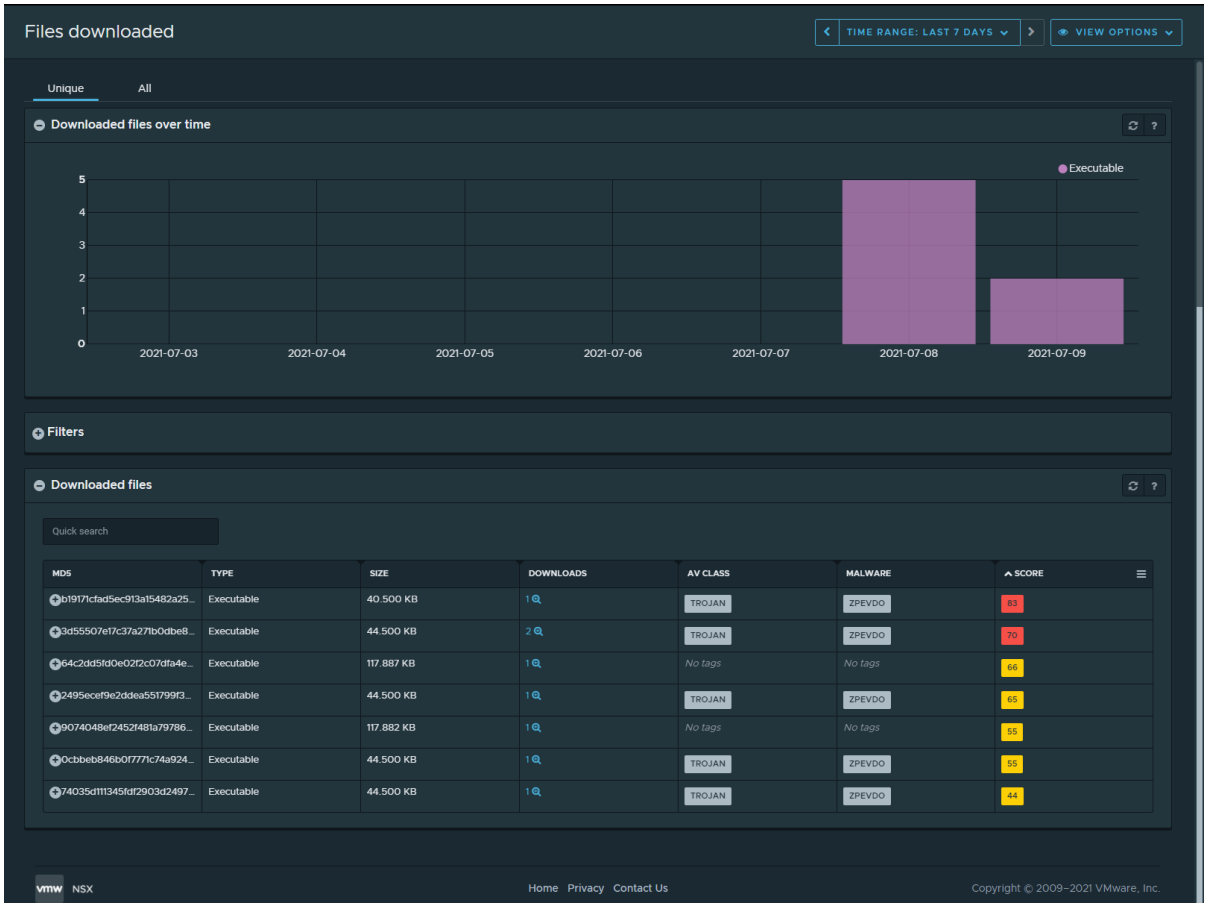
The page consists of several widgets that can be managed using the information in Getting Familiar with the NSX Network Detection and Response User Interface.

The page provides a high-level view of the numbers of files of different types that were downloaded in the network. It also allows you to look into the details of individual downloads, including access to the full reports of the analysis performed.

The following tabs are displayed on this page.

- The **Unique** tab displays distinct file downloads in the network that have been analyzed.
- The **All** tab displays all instances of file downloads that NSX Network Detection and Response analyzed in the network. Some of the files displayed are repetitions.

The following image shows a sample of the **Files downloaded** page.



Unique Tab

The **Unique** tab in the **Files downloaded** page of NSX Network Detection and Response displays distinct file downloads in the network that have been analyzed.

Downloaded Files Over Time

The **Downloaded files** widget provides an overview of the number of files that were downloaded in the monitored network during the specified time range. The graph is a daily histogram of downloaded files, grouped by the high-level file type.

The widget shows only distinct file downloads that have been analyzed.

The following are the displayed file types.

File Type	Description
Archive	Archive formats, such as ZIP or RAR.
Document	Includes other types of Office documents.
Executable	Binary application formats, such as Windows Portable Executable.
Java	Java application or applet.
Media	Macromedia (Adobe) Flash file.
Other	Other recognized file format.
PDF	Portable Document Format files.

File Type	Description
Script	An executable script, such as JavaScript, Python, and others.
Unknown	Unknown file type.

Use Filters in the Files Downloaded Page

NSX Network Detection and Response provides a filtering mechanism that allows you to focus on specific information about downloaded files that are of interest to you. The use of filters is optional.

Procedure

1 From the **Files Downloaded** page, click  to expand the **Filters** widget.

2 Click anywhere in the **Filter on** text box and select an item from the drop-down menu.

You can select from the following available filters. To further narrow the focus of the displayed information, you can combine multiple filters.

Filter Name	Description
Analysis tags	Restrict displayed files by their analysis tags. These are labels assigned to a file or URL by the system analysis. They can identify a threat or threat class, or refer to specific malicious behavior that was detected.
Analyst UUID	Restrict displayed files to the system analysis UUID for the downloaded file. This is an internal unique identifier for the analysis of a file.
Application protocol	Restrict displayed files transferred over one of the specified protocols. Supported values are HTTP/HTTPS, FTP, and SMB.
Contacted IP	Restrict displayed files to the IP address from which the file was downloaded. Like the Host IP filter, this supports IP addresses, CIDR blocks or IP address ranges.
File type filter	Restrict displayed files to one or more high-level file types. See the list of file types (above).
Files	Select Malicious to restrict displayed files to malicious files. These are files that were assigned a score of 70 or more (out of 100) by the system analysis.
Host IP	Restrict displayed files to the IP address of the host in the network that downloaded the file. This filter supports selecting one or more IP addresses, CIDR blocks (for example, 192.168.0.0/24) or IP address ranges (for example, 192.168.1.5-192.168.1.9).
HTTP Host	Restrict displayed files to the host name(s) from which the file was downloaded. **Note** This value is extracted from the HTTP Host header in the HTTP request that downloaded the file. Therefore, it is under the control of the client and can be spoofed by a malicious software, such as a malware binary already running on an infected host.
MD5	Restrict displayed files to the MD5 hash of the downloaded file.
Minimum score	Restrict displayed files to those assigned a score greater than your chosen value (from 1-100) by the system analysis.

3 To apply the selected filters, click **Apply**.

4 (Optional) To delete an individual filter, click the **REMOVE–** button next to its entry. To delete all the selected filters, click the **X** icon located on the right side of the **Filters** widget.

The **Filters** widget collapses when you delete all the selected filters.

Unique Downloaded Files List

The **Downloaded files** list displays all of the distinct files that have been downloaded by hosts in the network and processed by the NSX Advanced Threat Prevention service.

The **Quick search** text box in the upper-left corner of the list provides fast, as-you-enter search capability. It filters the rows in the list and displays only those rows that have text, in any column, that matches the query string that you entered in the search text box.

To customize the columns displayed in the list, click the ☰ icon located in the upper-right corner of the list.

You can customize the number of rows to be displayed. The default is 20 entries. Use the ‹ and › icons to navigate through multiple pages.

Each row is a summary of a downloaded file. Click the ⊕ icon or anywhere on an entry row to access a detailed view of the downloaded file.

The list is sorted by score and includes the following columns.

Column Name	Description
MD5	The MD5 hash of the downloaded file.
Type	The high-level file type of the downloaded file. Supported types are currently: ■ Archive – Archive formats such as ZIP or RAR ■ Document – Includes other types of Office documents ■ Executable – Binary application formats, such as Windows Portable Executable ■ Java – Java application or applet ■ Media – Macromedia (Adobe) Flash file ■ Other – Other recognized file format ■ PDF – Portable Document Format files ■ Script – An executable script such as JavaScript, Python, and others ■ Unknown – Unknown file type
Size	Size in bytes of the downloaded file.
Downloads	Number of times that the file was downloaded by hosts in the network. The displayed number and 🔍 provide a link to the detailed downloads page. The link passes an Analyst UUID filter that restricts the view to downloads of this specific file.
AV Class	A label defining the antivirus class of the downloaded file. If the label has a 🏷, you can click that icon for a description in a pop-up window.

Column Name	Description
Malware	A label defining the malware type of the downloaded file. If the label has a , you can click that icon for a description in a pop-up window.
Score	The score assigned to the downloaded file by the analysis indicates the critical level of the detected threat and ranges from 0-100: ■ Threats that are 70 or higher are considered to be critical. ■ Threats that are between 30-69 are considered to be medium-risk. ■ Threats that are between 1-29 are considered to be benign. For details about maliciousness core and risk estimate, see Analysis Report: Overview Tab. If the icon appears, it indicates the artifact has been blocked. The list is sorted by decreasing order (most critical threats at the top). Click to sort the list in increasing order (least critical threats at the top), then click to toggle back to the default.

Downloaded Files Details

The downloaded files details view is expanded within the **Downloaded Files** list.

You see a subset of the following available details, depending on which tab you have selected on the **Files Downloaded** page.

Detail Name	Description
Analysis report	Click the link or the icon to view the analysis report in a new tab.
File type	The high-level type of the downloaded file. See Downloaded Files Over Time for the list of file types.
File type details	If available, more details about the file type. For example, `PE executable, application, 32-bit, Intel i386` or `Zip archive data`.
Filename	If available, the name of the file.
Downloaded	For Unique downloads, the number of times that the file was downloaded by hosts in the network. Click the number or icon to view the file downloads on the downloads page. The link passes an Analyst UUID filter that restricts the view to downloads of the specific file.
Downloaded by	The IP address(es) of the host(s) in the network that downloaded the file. If available, click to view registration information and other data about the host in the WHOIS Pop-Up Window.
URL	The URL of the file download. This as a UTF-8 encoded Unicode string.
URL	The raw URL of the file download. If there are any non-ASCII characters in the URL, those, as well as the backslash character itself, will be backslash-encoded.
Protocol	Network protocols used to download the file. One of HTTP/HTTPS, FTP, or SMB.
Downloaded from	IP address of the contacted host. If available, click to view registration information and other data about the host in the WHOIS Pop-Up Windows.

Detail Name	Description
HTTP host	If available, the domain name of the contacted host. This name may be derived from other data including the IP address. If available, click to view registration information and other data about the host in the WHOIS Pop-Up Window.
User agent	The user agent string extracted from the HTTP/HTTPS request.
First download	For unique downloads, the timestamp of the first recorded detection of the file download.
Last download	For unique downloads, the timestamp of the most recent detection of the file download.
Timestamp	The timestamp of the detection of the file download.
File size	Size of the file in Bytes.
MD5	The MD5 hash of the downloaded file.
SHA1	The SHA1 hash of the downloaded file.
Submission status	Indicates why the downloaded file was not submitted for full analysis. Typically this is due to pre-filtering or other reasons. Hover your mouse over the icon to display a pop-up with further details.
Analyst UUID	The unique identifier returned by the NSX Advanced Threat Prevention service after processing the downloaded file.
Event ID	A link to the associated event for the file download. Click the ID or to view the event. See Detection Events for more information.

Analysis Overview

The analysis overview section provides a summary of the results of the analysis of a downloaded file by the NSX Advanced Threat Prevention service.

To open the full Analysis report in a new tab, click . See Using the Analysis Report.

To download the detected file to your local machine, click on the right side of the screen. From the drop-down menu, select **Download file** or **Download as ZIP**.

If you select **Download as ZIP**, the **Download file as a Zip** pop-up window appears, prompting you to provide an optional password for the archive. Click **Download** to complete downloading the `.ZIP` file.

Important The NSX Network Detection and Response application only allows you to download detected files under certain conditions.

If the artifact is considered low risk, is displayed and you can download it to your local machine.

If the artifact is considered risky, is not displayed unless your license has the `ALLOW_RISKY_ARTIFACT_DOWNLOADS` capability.

You must be aware that the artifact can possibly cause harm when opened.

The NSX Network Detection and Response interface might display the **Warning: Downloading Malicious File** pop-up window. Click the **I agree** button to accept the conditions and download the file.

For malicious artifacts, you might want to encapsulate the file in a `ZIP` archive to prevent other solutions that are monitoring your traffic from automatically inspecting the threat.

If you do not have the `ALLOW_RISKY_ARTIFACT_DOWNLOADS` capability and require the ability to download malicious artifacts, contact VMware Support.

Click and to expand and collapse the sections on the tab.

This Analysis Overview section provides a summary of the analysis results of a file or URL analyzed by the NSX Advanced Threat Prevention service. The section displays the following data.

- MD5 – The MD5 hash of the file. To search for other instances of this artifact in your network, click <search icon>.
- SHA1 – The SHA1 hash of the file.
- SHA256 – The SHA256 hash of the file.
- MIME Type – The label used to identify the type of data in the file.
- Submission – The submission timestamp

The Threat Level section starts with a summary of the analysis findings: `The file md5 hash was found to be malicious/benign.`

It then displays the following data:

Risk assessment

This section displays the risk assessment findings.

- Maliciousness score – Sets a score out of 100.

- Risk estimate – An estimate of the risk posed by this artifact:
  - High – This artifact represents a critical risk and you must address it in priority. Such subjects are typically Trojan files or documents that contain exploits, leading to major compromises of the infected system. The risks are multiple: from information leakage to system dysfunction. These risks are partially inferred from the type of activity detected. The score threshold for this category is usually greater than 70.
  - Medium – This artifact represents a long-term risk and you must monitor it closely. It can be a Web page containing suspicious content, potentially leading to drive-by attempts. It can also be an adware or a fake antivirus product that does not pose an immediate serious threat but can cause issues with the functioning of the system. The score threshold for this category is usually from 30-70.
  - Low – This artifact is considered benign and you can ignore it. The score threshold for this category is usually below 30.
- Antivirus class – The antivirus or malware class to which the artifact belongs. For example, a Trojan horse, worm, adware, ransomware, spyware, and so on.
- Antivirus family – The antivirus or malware family to which the artifact belongs. For example, valyria, darkside, and so on. To search for other instances of this family, click the search icon.

Analysis overview

The information displayed is sorted by severity and includes the following properties:

- Severity – A score between 0-100 of the maliciousness of the activities detected during analysis of the artifact. The additional icons indicate the operating systems that can run the artifact.
- Type – The types of activities detected during analysis of the artifact. These types include:
  - Autostart – Ability to restart after a machine shutdown.
  - Disable – Ability to disable critical components of the system.
  - Evasion – Ability to evade analysis environment.
  - File – Suspicious activity over the file system.
  - Memory – Suspicious activity within the system memory.
  - Network – Suspicious activity at the network level.
  - Reputation – Known source or signed by reputable organization.
  - Settings – Ability to permanently alter critical system settings.
  - Signature – Malicious subject identification.
  - Steal – Ability to access and potentially leak sensitive information.
  - Stealth – Ability to remain unnoticed by users.

    - Silenced – Benign subject identification.
- Description – A description corresponding to each type of activity detected during analysis of the artifact.
- ATT&CK Tactics – The MITRE ATT&CK stage or stages of an attack. Multiple tactics are separated by commas.
- ATT&CK Techniques – The observed actions or tools a malicious actor might use. Multiple techniques are separated by commas.
- Links – To search for other instances of this activity, click the search icon.

Additional artifacts

This section lists additional artifacts (files and URLs) that were observed during the analysis of the submitted sample and that were in turn submitted for in-depth analysis. This section includes the following properties:

- Description – Describes the additional artifact.
- SHA1 – The SHA1 hash of the additional artifact.
- Content type – The MIME type of the additional artifact.
- Score – The maliciousness score of the additional artifact. To view the associated analysis report, click  .

Decoded command line arguments

If any PowerShell scripts were executed during the analysis, the system decodes these scripts, making its arguments available in a more human-readable form.

Third-party tools

A link to a report on the artifact on VirusTotal portal.

All Tab

The **All** tab displays all instances of file downloads that were analyzed in your NSX network.

Downloaded Files Over Time in the All Tab

The **Downloaded files** widget in the **All** tab provides an overview of the number of files that were downloaded in the monitored network during the specified time range. The graph is a daily histogram of downloaded files, grouped by the high-level file type.

The widget shows all file downloads that have been analyzed.

See Downloaded Files Over Time for the list of file types.

Use Filters in the Files Downloaded Page

NSX Network Detection and Response provides a filtering mechanism that allows you to focus on specific information about downloaded files that are of interest to you. The use of filters is optional.

Procedure

1 From the **Files Downloaded** page, click  to expand the **Filters** widget.

2 Click anywhere in the **Filter on** text box and select an item from the drop-down menu.

You can select from the following available filters. To further narrow the focus of the displayed information, you can combine multiple filters.

Filter Name	Description
Analysis tags	Restrict displayed files by their analysis tags. These are labels assigned to a file or URL by the system analysis. They can identify a threat or threat class, or refer to specific malicious behavior that was detected.
Analyst UUID	Restrict displayed files to the system analysis UUID for the downloaded file. This is an internal unique identifier for the analysis of a file.
Application protocol	Restrict displayed files transferred over one of the specified protocols. Supported values are HTTP/HTTPS, FTP, and SMB.
Contacted IP	Restrict displayed files to the IP address from which the file was downloaded. Like the Host IP filter, this supports IP addresses, CIDR blocks or IP address ranges.
File type filter	Restrict displayed files to one or more high-level file types. See the list of file types (above).
Files	Select Malicious to restrict displayed files to malicious files. These are files that were assigned a score of 70 or more (out of 100) by the system analysis.
Host IP	Restrict displayed files to the IP address of the host in the network that downloaded the file. This filter supports selecting one or more IP addresses, CIDR blocks (for example, 192.168.0.0/24) or IP address ranges (for example, 192.168.1.5-192.168.1.9).
HTTP Host	Restrict displayed files to the host name(s) from which the file was downloaded. **Note** This value is extracted from the HTTP Host header in the HTTP request that downloaded the file. Therefore, it is under the control of the client and can be spoofed by a malicious software, such as a malware binary already running on an infected host.
MD5	Restrict displayed files to the MD5 hash of the downloaded file.
Minimum score	Restrict displayed files to those assigned a score greater than your chosen value (from 1-100) by the system analysis.

3 To apply the selected filters, click **Apply**.

4 (Optional) To delete an individual filter, click the **REMOVE–** button next to its entry. To delete all the selected filters, click the **X** icon located on the right side of the **Filters** widget.

The **Filters** widget collapses when you delete all the selected filters.

Downloaded Files List in the All Tab

The **Downloaded files** list displays all of the files that have been downloaded by hosts in the network and processed by the NSX Advanced Threat Prevention service.

The **Quick search** text box in the upper-left corner of the list provides fast, as-you-enter search capability. It filters the rows in the list and displays only those rows that have text, in any column, that matches the query string that you entered in the search text box.

To customize the columns displayed in the list, click the ☰ icon located in the upper-right corner of the list.

You can customize the number of rows to be displayed. The default is 20 entries. Use the ‹ and › icons to navigate through multiple pages.

Each row is a summary of a downloaded file. Click the ⊕ icon or anywhere on an entry row to access a detailed view of the downloaded file.

See Downloaded Files Details for additional information on the detailed view of the downloaded fiel.

The list is sorted by the timestamp information and includes the following columns.

Column Name	Description
Timestamp	The timestamp of the detection of the file download.
Host	The host that downloaded the file.
Contacted IP	IP address of the contacted host.
Location	For a download, this is the URL of the file in the supported format. For example, `\\127.0.0.2\samba_share\1128dedb.exe` for an SMB download or `http://www.example.com/download/example.zip` for an HTTP download. For an upload, "Upload" is displayed.
MD5	The MD5 hash of the downloaded file.
Type	The high-level type of the downloaded file. See the Downloaded Files Over Time for the list of file types.
AV Class	A label defining the antivirus class of the downloaded file. If the label has the 🏷 icon, you can click that for a pop-up description.
Malware	A label defining the malware type of the downloaded file. If the label has the 🏷 icon, you can click that for a pop-up description.
Score	The score assigned to the downloaded file by the NSX Intelligence analysis. Click ⇕ to sort the list by score. If ⊘ appears, it indicates the artifact has been blocked.

Using the Alert Management Page

The **Alert Management** page displays the rules for managing alerts in NSX Network Detection and Response.

NSX Network Detection and Response matches the events against the user-defined filters contained in these rules. Matching events are converted to `INFO` events (Demote) in the NSX Network Detection and Response UI, are deleted, or are assigned a custom impact value based on the selected action.

The **Custom rules** list defines the alert rules.

The quick search text box above the list provides the as-you-enter search feature. It filters the rows in the list, displaying only those rows that have text, in any column, that matches the query string.

Click  on the right side of the page to add a new alert rule. The **Manage alert** sidebar is displayed. See Working with the Manage Alert Sidebar for details.

You can customize the number of rows to be displayed. The default is 25 entries. To navigate through multiple pages, use the pagination icons.

The list is sorted by the Last Modified column and includes the following information.

Column Name	Description
Rule Name	The name of the alert rule. To sort the list by rule name, click  in the list header .
Expression	The matching expression of the rule is a number of filters that are matched against events. The expression may be truncated if it is too long. Expand the row to display the full content of the rule by clicking  or anywhere on the entry row. To sort the list by expression, click  in the list header.
Rule Action	The rule action defines what to do with an event that matches the expression: `demote` the event to `INFO`, `suppress` the event, or assign a custom `impact` value from 1 to 100. The action may be truncated if it is too long. Expand the row to display the full content of the rule by clicking the icon (or anywhere on the entry row). The rule name is appended to the action as a custom tag, for example `tag:network_event=rule_name`. To sort the list by rule action, click  in the list header.
Last Modified	The date and time of the last modification of the rule.
Actions	To view/edit the rule, click  . The **Manage alert** sidebar displays to allow you to view or make changes to the rule. To remove the rule, click  .

Working with the Manage Alert Sidebar

The **Manage Alert** sidebar allows you to create a rule that is matched against all subsequent events detected by NSX Network Detection and Response. When an event matches a rule, the rule action is applied.

Accessing the sidebar

You can access the **Manage Alert** sidebar in one of the following ways.

- From any tab on the **Host profile** page, click the **Host actions** button then select Manage alert from the pull-down menu. The sidebar panel is then prepopulated with relevant filters. You can edit these entries.
- Click the **Threats** tab on the **Host profile** page. On a threat card, click **Next steps** and select **Manage alert** from the pull-down menu.
- From the **Incident details** view, select a specific incident and click **Manage Alert**.

- From the **Alert Management** page, click  in the **Custom Rules** widget,

The **Manage alert** sidebar consists of three separate panels: FILTERS, ACTIONS, and REVIEW RULE. Each panel is displayed depending on which step of the Create Rule or Edit rule you are currently in.

You can close the **Manage alert** sidebar by clicking **X** in the upper-right corner. If you made changes, you must confirm the closing of the sidebar.

To create or edit a rule, you must perform three steps in the **Manage Alert** sidebar.

Step 1: Create or Edit Filters

The **Filters** tab has two edit modes that you can use when working with filters: Basic (the default) and Advanced. You can create or edit filters in either mode.

- To toggle the Create/Edit mode to Advanced mode, click the **Advanced** tab at the top of the sidebar.
- To toggle back to the Basic mode, click the **Basic** tab (but see the Important note).

To create a filter in Basic mode, perform the following steps.

1 Click **Add a new filter+**.

2 Select a filter from the filter entries drop-down menu.

The filters are grouped into four categories: Source, URL, Detection, and File. See the Attributes entries section in Alert Rule Syntax for more details about these categories.

3 Depending on the rule type selected, set its value. This may involve clicking a toggle, entering a value, selecting an item from a pull-down menu, or others.

To edit the filters, scroll through the list, select a filter, and modify the appropriate values. Delete an unwanted filter by clicking. You can also select more filters.

To create filters in Advanced mode, fill in the **Matching expression** text box, and add or edit a filter using the alert rules syntax. For example,

```
(network_event.relevant_host_ip: 10.154.115.91 OR network_event.relevant_host_ip:
10.1.1.1-10.255.255.255) AND NOT
(network_event.server_port: 53 OR network_event.server_port: 65535) OR
(network_event.other_host_hostname: block.lastline.com) AND
(network_event.threat: Lastline blocking test)
```

Important Normally you can toggle between the two sidebar edit modes, however if the matching expression filter you created or edited is not supported by the Basic mode, the **Basic** link is disabled and the **FILTERS** tab defaults to the Advanced editor.

Step 2: Define the action

After you define or edit a filter, to define the rule actions, click **Define Actions** in the bottom-right corner. The **Actions** panel has two edit modes: Basic actions (the default) and Advanced actions:

- Click the **Advanced actions** tab at the top of the sidebar to toggle the create/edit mode to Advanced mode.
- Click the **Basic actions** link to toggle back to the Basic mode.

There are two toggles on the **Actions** panel in Basic actions mode: **Manage alert** and **Custom impact (1-100)**.

Suppress action

1. Click the **Manage alert** toggle.
2. Select **Demote to INFO event** (the default) or **Delete** from the drop-down menu.

   The Demote action converts subsequent network events that match the rule into `INFO` events. Note that you must select INFO with the Event outcome filter.

   The Delete action deletes the matching events from the User Portal.

   Warning Any event that is deleted can no longer be accessed.

Custom impact

1. Click the **Custom impact (1-100)** toggle.
2. Click the radio buttons to select **Defined range** or **Single value**. If you selected **Defined range**, enter minimum and maximum values in the respective textboxes. If you selected **Single value**, enter the value in the textbox.

You can also define the actions using the Advanced actions panel.

1. Click the **Advanced actions** tab.
2. In the textbox, add or edit an action using the alert rules syntax.

   For example:

   ```
   demote:outcome=TEST
   ```

   or

   ```
   impact:min_impact=12,impact:max_impact=22
   ```

After you have selected the action, click **Review Rule** to go to the next step.

To correct the selected filters, click **Filters** to go back to the previous **Filters** panel.

Step 3: Review Rule

The Review Rule panel allows you to verify your alert rule.

1. In the Rule name text box, enter a name.

If you are editing an existing rule, you cannot change the name.

2 (Optional) Use the drop-down menu to select a license.

This drop-down menu is disabled if you launched the **Manage Alert** sidebar from the **Alert Management** page or if you are editing an existing rule.

3 In the **Rule summary** section, verify the selected filters that are listed.

If the **Filters** tab was left in Basic mode, the summary consists of a list of the selected filters. Each filter is displayed with its name and values. For example:

```
Rule summary
SERVER IP
12.6.6.6/32
RELEVANT HOST SILENCED
1
THREAT(S)
Torn rat
THREAT CLASS
Malicious file execution
```

If the **Filters** tab was left in Advanced mode, the summary displays the matching expression. For example:

```
Rule summary
(network_event.server_ip: 12.6.6.6/32) AND
(network_event.relevant_host_whitelisted: 1)
AND (network_event.threat: Torn RAT) AND
(network_event.threat_class: Malicious File
Execution)
```

If the **Actions** tab was left in Basic actions mode, the summary displays the action. For example:

```
SUPPRESSION ALERT
Demote to INFO event
```

If the **Actions** tab was left in Advanced actions mode, the summary displays the action. For example:

```
ACTION
impact:min_impact=12,impact:max_impact=22
```

4 (Optional) To correct the selected rule types, click **Edit rule** to go back to the previous page.

5 When you are done, click **Create Rule** to complete the rule or click **Update Rule** if you are editing an existing rule.

Alert Rule Syntax

You use the alert rule syntax to define the actions that NSX Network Detection and Response must take when events match a filter.

An alert rule consists of two parts: Matching expression and Actions.

Matching expression

A combination of clauses that express a condition on the attributes of an object.

A matching expression has the following format: `object_type . attribute_type: [relation]value`

The matching expression consists of the following four parts.

Part Name	Description
object_type	The object type to be matched. The following record type is supported: ■ network_event The object type and its attribute is separated by a dot (.).
attribute_type	The attribute to be matched (see Attribute entries). The object_type.attribute_type is separated from the [relation] and value by a colon (:).
[relation]	The relation between the object and its attribute and the value to match for. If no relation is specified, equality is the default. Supported relation types are: ■ Equality (:) ■ Greater than or equal (>, >=) ■ Less than or equal (<, <=)
value	The value to match against the object_type.attribute_type of the incoming events.

Multiple matching expressions are separated by the logical operators `AND`, `OR`, and `NOT`.

Actions

One or more modifications to be performed on the object.

An action has the following format: `action : target = value`

The action consists of three parts:

Part Name	Description
action	The action to be performed (see Supported actions). The action and its target are separated by a colon (:).
target	The supported target.
value	The optional value to apply to the target.

Multiple actions are separated by a comma (,) and are applied in the same order in which they were defined.

Attribute entries

The following list describes the different attribute entries that you can use when creating or updating new filters. The attributes are grouped into the following five categories.

SOURCE

Source Attribute	Description
client_ip	Matches an IP address or an IP address range. Address value must be an exact match. (network_event.client_ip: 142.42.1.6/24)
other_host_hostname	Matches the hostname of the other host associated with the event. Wildcard comparisons are supported: * for multiple characters, ? for single characters. You must escape (\) the wildcard characters to match a literal * or ?. (network_event.other_host_hostname: host.example.com)
other_host_in_homenet	If true, matches if the IP address of the other host associated with the event is in the home network. Expects a boolean value. (network_event.other_host_in_homenet: false)
other_host_ip	Matches an IP address or an IP address range. Address value must be an exact match. (network_event.other_host_ip: 10.10.4.2)
other_host_tag	Matches a host tag. Select an existing host tag. (network_event.other_host_tag: tag)
relevant_host_in_homenet	If true, matches if the IP address of the relevant host associated with the event is in the home network. Expects a boolean value. (network_event.relevant_host_in_homenet: true)
relevant_host_ip	Matches an IP address or an IP address range. Address value must be an exact match. (network_event.relevant_host_ip: 42.6.7.0/16)
relevant_host_tag	Matches a host tag. Select an existing host tag. (network_event.relevant_host_tag: tag)
relevant_host_whitelisted	Matches silenced source IP address. Expects a boolean value. (network_event.relevant_host_whitelisted: true)
server_ip	Matches an IP address or an IP address range. Address value must be an exact match. (network_event.server_ip: 12.6.6.6)
server_port	Matches a port number. Integer comparisons are performed: equality, inequality, greater-than, less-than, etc. (network_event.server_port: 7777)
transport_protocol	Matches either "TCP" or "UDP". (network_event.transport_protocol: UDP)

URL

URL Attribute	Description
full_url	Matches at least one URL in the event. Wildcard comparisons are supported: * for multiple characters, ? for single characters. You must escape (\) the wildcard characters to match a literal * or ?. For example, the query string character ? must be escaped (\?): (network_event.full_url: https://www.example.com/resource/path\?r=start&v=cK5G8fPmWeA)
normalized_url	Matches at least one normalized URL (a URL without the query string) in the event. Wildcard comparisons are supported: * for multiple characters, ? for single characters. You must escape (\) the wildcard characters to match a literal * or ?. (network_event.normalized_url: https://www.example.com/resource/path/)
resource_path	Matches at least one URL resource path in the event. Wildcard comparisons are supported: * for multiple characters, ? for single characters. You must escape (\) the wildcard characters to match a literal * or ?.

DETECTION

Detection Attribute	Description
custom_ids_rule_id	Matches an ID for an IDS rule. The numeric value must be an exact match. (network_event.custom_ids_rule_id: 987654321)
detector	Matches the name/unique identifier of the module that detected the event. The string value must be an exact match. (network_event.detector: llrules:1532130206460)
event_outcome	Matches either "DETECTION" or "INFO". (network_event.event_outcome: DETECTION)
event_type	Matches one of "BINARYDOWNLOAD", "DNS", "DNSANOMALY", "DYNAMICIP", "HTTPANOMALY", "IDS", "IP", "LLANTARULE", "NETFLOW", "NETFLOWANOMALY", "NETWORK", "TLSANOMALY", or "URL". (network_event.event_type: IDS)
llanta_rule_uuid	Matches the UUID of a system rule. The numeric value must be an exact match. (network_event.llanta_rule_uuid: b579caeec719415cb04f925f8f187cb0)
operation	Matches one of "BLOCK", "INFO", "LOG", or "TEST". (network_event.operation: BLOCK)
threat	Matches a valid string defining a threat. Wildcard comparisons are supported: * for multiple characters, ? for single characters. You must escape (\) the wildcard characters to match a literal * or ?. (network_event.threat: Torn RAT)
threat_class	Matches a threat class. The string value must be an exact match. (network_event.threat_class: Malicious File Execution)

FILE

File Attribute	Description
av_class	Matches at least one av_class analysis tag. The string value must be an exact match. (network_event.av_class: exploit)
file_category	Matches one of the supported categories of files. The string value must be an exact match. (network_event.file_category: Java)
file_md5	Matches a valid MD5 sum. (network_event.file_md5: bb4f64ddfb8704d2bf69b0216be7f837)
file_sha1	Matches a valid SHA1 sum. (network_event.file_sha1: c3e266ede7f6fec7a021a4ae0edf248848d5ae06)
file_size	Matches a file size in bytes. It must be a valid integer. Integer comparisons are performed: equality, inequality, greater-than, less-than, etc. (network_event.file_size: > 1042249837)
file_type	Matches a valid string defining a file type. Wildcard comparisons are supported: * for multiple characters, ? for single characters. You must escape (\) the wildcard characters to match a literal * or ?. (network_event.file_type: ?xecutable)
malware	Matches at least one av_family or lastline_malware analysis tag. The string value must be an exact match. (network_event.malware: emotet)
malware_activity	Matches at least one activity analysis tag. The string value must be an exact match. (network_event.malware_activity: Execution: Spawning Powershell with too many parameters)

OTHER

Other Attribute Name	Description
custom_tag	Matches a user-defined tag assigned to events. The string value must be an exact match. (network_event.custom_tag: tagged_event)

Supported actions

The following are the actions that you can use when defining rules.

Action Name	Description
demote	Demotes the outcome of the matching event to a different mode. Supported targets: outcome. Allowed values: "INFO" or "TEST".
impact	Set a lower or upper bound on the impact of an event. Supported targets: ■ impact: Sets the lower and upper bound to the same value. ■ max_impact: Sets the upper bound on impact. Less or equal to value. ■ min_impact: Sets the lower bound on impact. Greater or equal to value. Allowed values: an integer from 1-100.

Action Name	Description
`suppress`	Suppresses all threats on the matching event. This results in it being scored as a false positive with an impact of zero (0), which effectively deletes the event. Supported targets: `network_event`. Allowed values: none.
`tag`	Assign a user-defined tag to the matching event. Supported targets: `network_event`. Allowed values: a valid string.

Using the Analysis Report

The analysis report produced by NSX Network Detection and Response contains the detailed results of an analysis performed by the NSX Advanced Threat Prevention service on a submitted file.

Besides the maliciousness score, the report also contains important information about the activity of the analysis subject. The described activity constitutes the base of the NSX Network Detection and Response threat assessment and scoring.

The analysis report starts in the **Overview** tab.

Analysis Report: Overview Tab

The **Overview** tab in the **Analysis Report** page of the NSX Network Detection and Response UI provides a summary of the analysis results for the file analyzed by the NSX Advanced Threat Prevention service.

To download the detected file to your local machine, click  on the right side of the screen. From the drop-down menu, select **Download file** or **Download as ZIP**.

If you select **Download as ZIP**, the **Download file as a Zip** pop-up window appears, prompting you to provide an optional password for the archive. Click **Download** to complete downloading the `.ZIP` file.

Important The NSX Network Detection and Response application only allows you to download detected files under certain conditions.

If the artifact is considered low risk, is displayed and you can download it to your local machine.

If the artifact is considered risky, is not displayed unless your license has the `ALLOW_RISKY_ARTIFACT_DOWNLOADS` capability.

You must be aware that the artifact can possibly cause harm when opened.

The NSX Network Detection and Response interface might display the **Warning: Downloading Malicious File** pop-up window. Click the **I agree** button to accept the conditions and download the file.

For malicious artifacts, you might want to encapsulate the file in a `ZIP` archive to prevent other solutions that are monitoring your traffic from automatically inspecting the threat.

If you do not have the `ALLOW_RISKY_ARTIFACT_DOWNLOADS` capability and require the ability to download malicious artifacts, contact VMware Support.

Analysis Overview Section

Note If the NSX Advanced Threat Prevention service encountered errors during the file analysis, a highlighted block is displayed. It contains a list of the errors encountered.

This Analysis Overview section provides a summary of the analysis results of a file or URL analyzed by the NSX Advanced Threat Prevention service. The section displays the following data.

- MD5 – The MD5 hash of the file. To search for other instances of this artifact in your network, click <search icon>.
- SHA1 – The SHA1 hash of the file.
- SHA256 – The SHA256 hash of the file.
- MIME Type – The label used to identify the type of data in the file.
- Submission – The submission timestamp

Threat Level Section

The Threat Level section starts with a summary of the analysis findings: `The file md5 hash was found to be malicious/benign.`

It then displays the following data:

Risk assessment

This section displays the risk assessment findings.

- Maliciousness score – Sets a score out of 100.
- Risk estimate – An estimate of the risk posed by this artifact:
  - High – This artifact represents a critical risk and you must address it in priority. Such subjects are typically Trojan files or documents that contain exploits, leading to major compromises of the infected system. The risks are multiple: from information leakage to system dysfunction. These risks are partially inferred from the type of activity detected. The score threshold for this category is usually greater than 70.
  - Medium – This artifact represents a long-term risk and you must monitor it closely. It can be a Web page containing suspicious content, potentially leading to drive-by attempts. It can also be an adware or a fake antivirus product that does not pose an immediate serious threat but can cause issues with the functioning of the system. The score threshold for this category is usually from 30-70.
  - Low – This artifact is considered benign and you can ignore it. The score threshold for this category is usually below 30.
- Antivirus class – The antivirus or malware class to which the artifact belongs. For example, a Trojan horse, worm, adware, ransomware, spyware, and so on.
- Antivirus family – The antivirus or malware family to which the artifact belongs. For example, valyria, darkside, and so on. To search for other instances of this family, click the search icon.

Analysis overview

The information displayed is sorted by severity and includes the following properties:

- Severity – A score between 0-100 of the maliciousness of the activities detected during analysis of the artifact. The additional icons indicate the operating systems that can run the artifact.
- Type – The types of activities detected during analysis of the artifact. These types include:
  - Autostart – Ability to restart after a machine shutdown.
  - Disable – Ability to disable critical components of the system.
  - Evasion – Ability to evade analysis environment.
  - File – Suspicious activity over the file system.
  - Memory – Suspicious activity within the system memory.
  - Network – Suspicious activity at the network level.
  - Reputation – Known source or signed by reputable organization.
  - Settings – Ability to permanently alter critical system settings.
  - Signature – Malicious subject identification.

- Steal – Ability to access and potentially leak sensitive information.
- Stealth – Ability to remain unnoticed by users.
- Silenced – Benign subject identification.

- Description – A description corresponding to each type of activity detected during analysis of the artifact.
- ATT&CK Tactics – The MITRE ATT&CK stage or stages of an attack. Multiple tactics are separated by commas.
- ATT&CK Techniques – The observed actions or tools a malicious actor might use. Multiple techniques are separated by commas.
- Links – To search for other instances of this activity, click the search icon.

Additional artifacts

This section lists additional artifacts (files and URLs) that were observed during the analysis of the submitted sample and that were in turn submitted for in-depth analysis. This section includes the following properties:

- Description – Describes the additional artifact.
- SHA1 – The SHA1 hash of the additional artifact.
- Content type – The MIME type of the additional artifact.
- Score – The maliciousness score of the additional artifact. To view the associated analysis report, click .

Decoded command line arguments

If any PowerShell scripts were executed during the analysis, the system decodes these scripts, making its arguments available in a more human-readable form.

Third-party tools

A link to a report on the artifact on VirusTotal portal.

Analysis Report: Report Tab

The information displayed on the **Report** tab changes depending on the type of file that NSX Network Detection and Response analyzed.

To view a report, click the down-arrow on the **Report** tab and select one of the available reports.

Click and to expand and collapse the sections on the tab.

Analysis Information Section

The **Analysis Information** section contains key information about the analysis that the current report refers to:

- Analysis subject: The MD5 hash of the file.

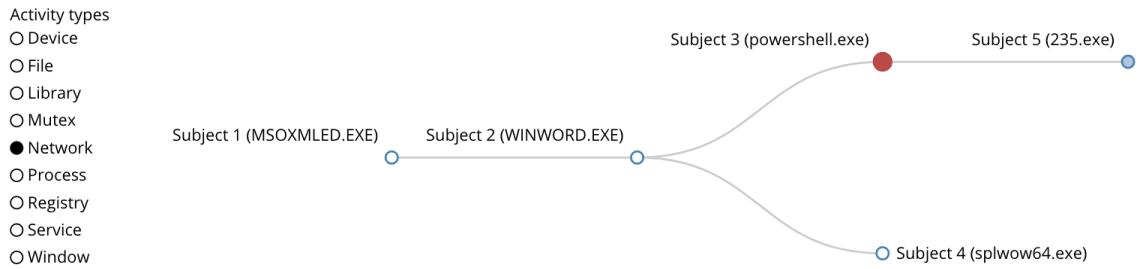
- Analysis type: The type of analysis that was performed:
  - Dynamic analysis on Microsoft Windows 10: The NSX Advanced Threat Prevention service ran the analysis subject in a simulated Windows 10 environment using the NSX Network Detection and Response sandbox. The system monitors the file behavior and its interactions with the operating system looking for suspicious or malicious indicators.
  - Dynamic analysis on Microsoft Windows 7: The NSX Advanced Threat Prevention service ran the analysis subject in a simulated Windows 7 environment using the NSX Network Detection and Response sandbox. The system monitors the file behavior and its interactions with the operating system, looking for suspicious or malicious indicators.
  - Dynamic analysis in instrumented Chrome browser: The NSX Advanced Threat Prevention service inspected the analysis subject (such as an HTML file or URL) using the instrumented browser, which is based on Google Chrome. The instrumented browser reproduces faithfully the behavior of the real browser and therefore is not easily fingerprinted by malicious content.
  - Dynamic analysis in emulated browser: The NSX Advanced Threat Prevention service inspected the analysis subject (such as an HTML file or URL) using the emulated browser. The emulated browser can dynamically emulate different browser "personalities" (for example, changing its `user-agen` or varying the APIs that it exposes). This capability is useful when analyzing malicious content that targets specific browser types or versions. The drawback of this type of analysis is that this browser is less realistic and can possibly be fingerprinted by malicious content.
  - Dynamic analysis in simulated file-viewer: The NSX Advanced Threat Prevention service inspected the analysis subject (such as a PDF file) using the simulated file-viewer. The viewer can detect embedded contents and links.
  - Archive inflation: The NSX Advanced Threat Prevention service inflated the analysis subject (an archive), extracted its contents, and submitted the contents for analysis if they are of an appropriate type.
- Password used: If available, the password that the NSX Advanced Threat Prevention service used to decrypt the sample successfully, is provided.

Analysis Relationships Widget

A single analysis might require the NSX Advanced Threat Prevention service to monitor multiple subjects.

For example, during a file analysis, the original application might launch multiple processes. Similarly, during a URL analysis, more URLs might be referenced and fetched.

In this case, NSX Network Detection and Response generates the **Analysis subjects overview** widget, which provides a graphical representation of the relationship of each analysis subject that the NSX Advanced Threat Prevention service monitored during the analysis.



The widget displays a node for each analysis subject. Two nodes are linked by an edge if the corresponding analysis subjects were found to interact during the analysis (for example, a process started another process).

On the left-hand side of the widget is a legend of activities that were observed during the analysis. Click the radio button next to an activity name to highlight the analysis subjects that displayed that specific activity. You can also select a set of activities.

Click a node to collapse the subsequent related nodes.

Double-clicking a node takes you to the section of the report that provides detailed information about the corresponding analysis subject.

Analysis File Report

The **Analysis subject** sections display detailed information about the file or files contained or accessed by the sample when the NSX Advanced Threat Prevention service processed it.

To expand the section, click ⊕.

For an executable file, the following data is displayed:

- Name: The name of the executable, if available.
- MD5: The MD5 hash of the file.
- SHA1: The SHA1 hash of the file.
- File type: The type of executable, for example, `PE executable, application, 32-bit, Intel i386`.
- File size: The file size.
- Command line: The full command line, including any arguments or options. For example, `C:\Users\ExampleUser\AppData\Local\Temp\exe_malware.exe`.
- Execution context: The privilege level invoked by the executable.
- Architecture: The architecture of the executable.
- Analysis reason: Why the started processing the file.

Analysis File Activities

The **Analysis subject** sections display the actual activity of the sample, as collected by the NSX Advanced Threat Prevention service.

The sections include the original subject being analyzed and additional subjects tracked by the analysis environment because they were either spawned by the original subject or because the original subject tampered with their memory.

Note Not all of these activities are present for a specific sample.

Click the  icon to expand each of the following sections.

Section Name	Description
Console I/O	Data written to console handles (standard input and standard output file descriptors).
Decoded command line arguments	The arguments to malicious PowerShell scripts are often encoded or obfuscated. If a script was executed during the analysis, the VMware backend decodes it, making its arguments available in a more human-readable form.
Device I/O	Device I/O List of I/O operations attempted by the subject during runtime. For each operation, the targeted device and the control code are recorded.
Driver activity	List of drivers accessed by the subject during runtime. The following operations are recorded: loading and unloading.
Exceptions	List of scripts executed by the subject during runtime. For each row, there is an entry for the Name, TYPE, and INTERPRETER. You can sort the list by any column.
Executed scripts	List of scripts executed by the subject during runtime. For each row, there is an entry for the Name, TYPE, and INTERPRETER. You can sort the list by any column.
File system activity	List of files accessed by the subject during runtime. The following operations are recorded: reading, writing, renaming, deletion. For written files, the new size and MD5 hash of the file is recorded.
Libraries	List of library files loaded by the subject during runtime.
Memory contents	Noteworthy data found in program memory. The system extracts, for example, IPs, domains, and URLs during analysis.
Mutex activity	List of mutex locks accessed by the subject during runtime. The following operations are recorded: creation and opening.
Network activity	List of network conversations involving the subject during runtime. The following type of conversations are recorded: communications over FTP, HTTP, IRC, SMTP, and other types of UDP/TCP protocols. DNS requests and remote file downloads are also recorded.
Process interactions	List of process interactions attempted by the subject during runtime. The following operations are recorded: process creation, thread creation, memory reading and writing.
Registry activity	List of registry keys and values accessed by the subject during runtime. The following operations are recorded: reading, writing, deletion and monitoring.
Service activity	List of services accessed by the subject during runtime. The following operations are recorded: starting, stopping, modifying parameters.
Windows activity	List of windows opened by the subject during runtime.

Analysis File Artifacts

The **Events report** section displays additional artifacts that the NSX Advanced Threat Prevention service gathers while it processes the sample.

These artifacts are included in the report for you to view.

Packet Capture

If the subject generated network traffic, this traffic is collected and displayed in the captured traffic widget.

Extracted Files

For an inflated archive, a list of the contents is displayed. Each row shows the mime type, tag (indicates the type of analysis), description, filename (if available from the archive), and score of the artifact. A score is provided only if the artifact is analyzed. In this case, a link to its report is also provided.

If the NSX Advanced Threat Prevention service encountered an error when unpacking an archive, it displays an alert indicating the error condition. Errors include maximum file limit exceeded, maximum depth limit exceeded, and maximum child task limit exceeded.

Generated Files

During analysis, the sample might generate various files. These files are displayed in a list sorted by PATH.

- PATH: The path of the artifact in the file system.
- TYPE: The determined file type. To sort the list by file type, click **TYPE**.

Click the  icon to expand a row. Data for MD5, SHA1, Size (bytes), Packers, and Signatures are displayed. Data might not be available for all fields.

Decoded Command Line Arguments

The arguments to malicious PowerShell scripts are often encoded or obfuscated. If a script was executed during the analysis, the NSX Advanced Threat Prevention service decodes it, making its arguments available in a more human-readable form. These arguments are displayed in a list showing the analysis subject and decoded script.

Analysis URL Report

The **Analysis details** section displays the actual activities of the analysis subject, as collected by the NSX Advanced Threat Prevention service. An activity is used to determine an assessment of its type.

The following activities are displayed in this **Analysis details** section.

Activity Type	Description
Network activity	Lists all URLs visited during the analysis, as well as additional web content requested or contained by the subject. Each additional URL is recorded together with its content type, the server status code, the server IP address, the response content hashes (MD5 and SHA1), the response content length, and the timing of the request (start time, end time, and duration in milliseconds).
Resources	Lists local resources that were accessed during the URL analysis via the res protocol. Malicious web pages sometimes access local resources to probe the execution environment; for example, to determine if certain programs are installed. This section is displayed only if resources events were encountered during analysis.

Activity Type	Description
Code execution activity	Lists code that was executed during the analysis. In particular, it displays interesting code that was statically included in a resource (using a `<script>` tag), and all the code that was dynamically generated and executed during the URL analysis. Malicious code is often generated at runtime in order to bypass static signatures and to make its analysis more complicated. ■ Static JavaScript code: Displayed only if relevant events were encountered during analysis. ■ Dynamic JavaScript code: Report indicates if no events were encountered during analysis. ■ HTML code: Code that has been added to the document dynamically through functions like `document.write()`. Report otherwise indicates if no events were encountered during analysis.
Hidden iframes	Lists hidden HTML tags, such as `iframe`, that have been detected during the navigation. Hidden elements are sometimes used in compromised pages to pull in malicious code from third-party websites. This section is displayed only if hidden tags were encountered during analysis.
Memory contents	Lists the strings that were found during the analysis. This section is displayed only if strings were encountered during analysis.
Textual content	Shows the textual content extracted from a document. This section is displayed only if text was found during analysis, PDF analysis only.
Links in documents	Shows the links that were found in analyzed documents. This section is displayed only if links were encountered during the analysis.
Plugins	Lists any use of common browser plugins. Calls to these plugins are recorded and the report contains the details about the invoked methods and the passed arguments.
Applets	Shows the Java applets that were downloaded during the URL analysis. This section is displayed only if applets were found during analysis.
Exploits	The analysis environment has the capability to detect shellcode contained in analysis subjects. Detected shellcode are extracted and included in the report in hexadecimal format.
Shellcode	The analysis environment has the capability to detect shellcode contained in analysis subjects. Detected shellcode are extracted and included in the report in hexadecimal format.
Processes	Lists the processes that were spawned during the URL analysis. This section is displayed only if spawned processes were found during analysis.
Dropped Files	Lists files that were stored on the system hard disk during the URL analysis. This section is displayed only if file operations were encountered during analysis.

Administering NSX Network Detection and Response

You can upgrade and delete the NSX Network Detection and Response feature using the NSX Application Platform page in the **System** tab.

Note The NSX Malware Prevention and NSX Intelligence features are each interdependent on the NSX Network Detection and Response feature.

When the NSX Network Detection and Response and the NSX Malware Prevention features are both activated, the NSX Malware Prevention feature publishes file events to NSX Network Detection and Response. When the NSX Network Detection and Response and the NSX Intelligence features are both activated, NSX Intelligence publishes anomalous events to NSX Network Detection and Response.

As a consequence, when you activate, update, or delete one of the interdependent features, the system reconfigures the other feature to update the features' interdependence status.

For example, if you delete the NSX Network Detection and Response feature when the NSX Malware Prevention feature is activated, the NSX Malware Prevention feature card briefly gives a status of `Deployment in progress` as the system reconfigures the feature to stop sending file events to the newly deleted NSX Network Detection and Response feature. Similarly, if you delete the NSX Malware Prevention feature when the NSX Network Detection and Response is activated, the NSX Network Detection and Response feature card briefly displays the `Deployment in progress` status while the system updates the interdependence status of the two features.

Upgrading NSX Network Detection and Response

Upgrading the NSX Network Detection and Response feature occurs only when you upgrade the NSX Application Platform.

See the "Upgrade the NSX Application Platform" topic in the *Deploying and Managing the VMware NSX Application Platform* document for versions 3.2 or later delivered in the VMware NSX Documentation set.

Delete NSX Network Detection and Response

If for some reason you must delete the NSX Network Detection and Response feature, use the steps described in this section.

Note When you delete the NSX Network Detection and Response feature, the system deletes all the data analytics that have been gathered, along with the feature if the NSX Malware Prevention feature is currently not activated. The action is permanent and results in loss of previously collected data.

If the NSX Malware Prevention feature is still activated, the data analytics that have been gathered remains intact when the NSX Network Detection and Response feature is deleted until both features are deleted.

Prerequisites

- NSX Application Platform must be in a good state and there are no active alarms.
- You must be logged in using an Enterprise Admin account.

Procedure

1. From your browser, log in with the required privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
2. In the NSX Manager UI, select **System > NSX Application Platform** in the Configuration section.
3. Navigate to the **Features** section and in the NSX Network Detection and Response feature card, click **Actions** and select **Delete**.
4. Click **Delete** in the **Delete NSX Network Detection and Response** dialog box.

Results

After a successful deletion, the NSX Network Detection and Response feature card returns to a ready-to-activate state.

Troubleshooting NSX Network Detection and Response

This section provides information about how to resolve some problems you might encounter when activating the NSX Network Detection and Response feature.

NSX Network Detection and Response Cloud Region Information Retrieval Failed

The NSX Network Detection and Response activation wizard failed to retrieve information about the available cloud regions.

Problem

NSX Network Detection and Response integrates with the NSX Advanced Threat Prevention cloud service for processing detection event data. Before you can activate the NSX Network Detection and Response feature, you must select a supported cloud region to which information about suspicious or malicious detection events is sent for processing. If the activation wizard fails to obtain the information about the available cloud regions, the feature activation is blocked from progressing.

Cause

The NSX Network Detection and Response activation service must connect to `nsx.lastline.com` on TCP port 443 to retrieve the list of available cloud regions. If the connection is unavailable, the activation wizard cannot retrieve this list.

Solution

1 If your NSX Manager appliance is configured to use a web proxy for Internet-bound connections, ensure that the web proxy is configured correctly and is reachable from the workloads running in the Kubernetes cluster used for NSX Application Platform.

2 Ensure that NSX Application Platform is deployed correctly and is reported as `STABLE` on the **Systems > NSX Application Platform** UI page.

3 Close the NSX Network Detection and Response wizard and refresh your web browser to force the UI to reinitialize the list of available cloud regions.

NSX Network Detection and Response Activation Precheck Failed

The NSX Network Detection and Response feature activation wizard reports an error while performing the precheck step.

Problem

The NSX Network Detection and Response activation wizard performs prechecks to verify the connectivity between the NSX Advanced Threat Prevention cloud service and the Kubernetes cluster running the NSX Application Platform. Any errors encountered are displayed on the NSX Network Detection and Response activation wizard. If the activation wizard reports an error during the precheck step, the NSX Network Detection and Response feature activation becomes blocked and the **Activate** button remains dimmed.

Following are some of several errors you might encounter if the activation precheck failed.

- `Cloud regions APIs returned invalid data (missing or invalid data).`
- `Contacting cloud API for validating the NSX license failed`
- `The NSX Installation does not have the required license.`

Cause

The precheck step validates the connectivity from the Kubernetes cluster running the NSX Application Platform to the NSX Advanced Threat Prevention cloud region that you selected. The precheck step also validates the available NSX Data Center licenses you are entitled to use for the NSX Network Detection and Response feature. If the connectivity precheck fails, the deployment wizard cannot validate the license eligibility.

Solution

1 If your NSX Manager appliance is configured to use a web proxy for Internet-bound connections, ensure that the web proxy is configured correctly and is reachable from the workloads running in the Kubernetes cluster used for NSX Application Platform.

2 Ensure that NSX Application Platform is deployed correctly and is reported as `STABLE` on the **Systems > NSX Application Platform** UI page.

3 The deployment precheck can take up to 30 minutes to deploy and validate NSX Advanced Threat Prevention cloud reachability and NSX Data Center license eligibility. Wait for the precheck items to complete and verify the outcome for each row.

   a For any item marked as `Failed`, point to the icon to view details.

   b Ensure the license requirements stated as part of a failure are satisfied.

   c If the error indicates connectivity errors, ensure that the NSX Application Platform can communicate with the Internet.

4 If the precheck failure does not provide information about the failure, gather additional information.

   a Collect an NSX Application Platform support bundle and inspect the logs for any Kubernetes pod with the name starting with `nsx-ndr-precheck`.

   b Alternatively, the logs can also be queried interactively on the NSX Manager appliance using the following steps.

      1 Log in to the NSX Manager appliance as **`root`**.

      2 Use the following command to mark the Kubernetes configuration for any subsequent `helm` and `kubectl` invocations.

```
export KUBECONFIG=/config/vmware/napps/.kube/config
```

      3 Using the following command, ensure that the NSX Network Detection and Response precheck helm chart was deployed successfully.

```
helm --namespace nsxi-platform list --all --filter 'nsx-ndr-precheck'
```

      Verify that the `STATUS` property displays `deployed`.

      4 Using the following command, inspect the events for any precheck pods that have been deployed.

```
kubectl --namespace nsxi-platform describe pod --selector='app.kubernetes.io/
instance=nsx-ndr-precheck'
```

      The `Events` section provides status of the precheck jobs and any actions associated with those jobs.

      5 To inspect the logs for any precheck pods that have been deployed, use the following command.

```
kubectl --namespace nsxi-platform get pods --selector='app.kubernetes.io/
instance=nsx-ndr-precheck' -o wide
```

      For each pod listed in `RUNNING` or `COMPLETED` state, view the logs using the following command.

```
kubectl --namespace nsxi-platform logs --container=main <pod-name>
```

5 After resolving the reported errors, try activating the NSX Network Detection and Response feature again.

Certificate Verification Error While Connecting to the NSX Manager

Precheck fails while activating NSX Network Detection and Response.

Problem

The logs of the `nsx-ndr-precheck` job contain the following error:

```
2022-10-04 19:43:44,954 - nsx_api_client.nsx_api_client - ERROR - communication
error: HTTPSConnectionPool(host='external-nsx-manager', port=443): Max retries
exceeded with url: /api/v1/licenses (Caused by SSLError(SSLError(1, '[SSL:
CERTIFICATE_VERIFY_FAILED] certificate verify failed (_ssl.c:852)'),))
```

In these logs, the specific issue manifests with the error string `"CERTIFICATE_VERIFY_FAILED"` in the connection to the host `"external-nsx-manager"`.

To know more information about how to access these logs, refer step 4 b of NSX Network Detection and Response Activation Precheck Failed.

Cause

The problem occurs because of issues with the certificate setup for the NSX Manager cluster.

Solution

1. To resolve the issue, make sure that you have correctly configured the certificates for each node in the NSX Manager cluster and for the cluster itself. Specifically verify:
   - Each certificate has a common name that matches the fully qualified domain name it is used for.
   - No two certificates use the same common name.

2 To verify that the certificate setup is correct, here are the two suggested approaches.

a Using a web browser:

With a web browser, visit each of the NSX Manager nodes individually using their fully qualified domain names, and also visit the cluster domain name. At each of these domains, use the browser tools and view the certificate information and in particular **Common Name** of the certificate. Confirm that the common name exactly matches the visited domain name in the URL bar of the browser.

b Using the Command Line Tools:

In a shell, for one of the NSX Manager nodes, run the following commands:

- Extract the trusted CA certificate from the Kubernetes platform.

```
napp-k get secret/nsx-ndr-precheck-ccpchk-nsx-manager-truststore   -o
                  jsonpath='{.data.ca\.crt}'| base64 -d > /tmp/ca.crt
```

The command extracts the trusted certificate authorities from the Kubernetes platform into the `/tmp/ca.crt` file.

- View the `subject` of the trusted certificates.

```
openssl storeutl -noout -certs --text /tmp/ca.crt |grep "Subject:"
       Subject: C=US, ST=CA, L=Palo Alto,
                    O=VMware Inc., OU=NSX,   CN=k8s-platform-ca
       Subject: C=US, ST=CA, L=Palo Alto,
                    O=VMware Inc., OU=NSX,
                CN=vmnsxt-mgmt-p01.example.com
       Subject: C=US, ST=CA, L=Palo Alto,
                    O=VMware Inc., OU=NSX,
                CN=vmnsxt-mgmt-p02.example.com
       Subject: C=US, ST=CA, L=Palo Alto,
                    O=VMware Inc., OU=NSX,
                CN=vmnsxt-mgmt-p03.example.com
       Subject: C=US, ST=CA, L=Palo Alto,
                    O=VMware Inc., OU=NSX,
                CN=vmnsxt-mgmt-p01.example.com
```

The command parses those certificates and prints out the `subject` of each certificate. The command also includes the common name of the certificate which is after `"CN="` in the sample example.

Verify that the common names shown are all different and match the expected fully qualified domain names for the nodes of the NSX Manager cluster and for the cluster itself.

In the earlier example, the common name `vmnsxt-mgmt-p01.example.com` occurs twice. The sample setup is configured incorrectly. The fully qualified domain name for the NSX Manager cluster is `vmnsxt-mgmt.example.com`, but the certificate uses an incorrect common name `vmnsxt-mgmt-p01.example.com`.

NSX Cloud Connector Deployment Failed

The NSX Network Detection and Response activation wizard reports an error after attempting to deploy the NSX Cloud Connector component.

Problem

As part of the NSX Network Detection and Response feature activation, the activation wizard attempts to deploy the NSX Cloud Connector component. The NSX Cloud Connector registers the NSX installation with the NSX Advanced Threat Prevention cloud service and builds a secure channel between the local resources and the cloud resources. If an issue is encountered during those steps, the activation wizard reports an NSX Cloud Connector deployment error, which blocks the NSX Network Detection and Response activation to continue and it eventually times out.

Cause

The NSX Cloud Connector establishes connectivity to the NSX Advanced Threat Prevention cloud service region that you selected and it triggers registration using the NSX licenses. If the connection is unavailable, the registration fails or times out.

Solution

1. The NSX Cloud Connector uses the same communication channel that was previously validated during the NSX Network Detection and Response activation precheck. If the NSX configuration changed between running the NSX Network Detection and Response precheck and the actual NSX Network Detection and Response activation, rerun the activation precheck. If you encounter any error, follow the troubleshooting information for activation precheck failure.
2. Ensure that NSX Application Platform is deployed correctly and is reported as `STABLE` on the **Systems > NSX Application Platform** UI page.

3 Inspect the logs for the NSX Cloud Connector registration service.

   a Collect an NSX Application Platform support bundle and inspect the logs for any Kubernetes pod with the name starting with `cloud-connector-register`.

   b Alternatively, the logs can also be queried interactively on the NSX Manager appliance using the following steps.

      1 Log into the NSX Manager appliance as **`root`**.

      2 Use the following command to mark the Kubernetes configuration for any subsequent `helm` and `kubectl` invocations.

```
export KUBECONFIG=/config/vmware/napps/.kube/config
```

      3 Using the following command, ensure that the NSX Cloud Connector helm chart is deployed successfully.

```
helm --namespace nsxi-platform list --all --filter 'cloud-connector'
```

      Verify that the `STATUS` property displays `deployed`.

      4 Inspect that the registration pod is deployed and completed successfully.

```
kubectl --namespace nsxi-platform get pods --selector='job-name=cloud-connector-
register'
```

      The pod should show the `STATUS` as `Completed`.

      5 Inspect events for the registration pod, using the following command.

```
kubectl --namespace nsxi-platform describe pod --selector='job-name=cloud-
connector-register'
```

      The Events section provides status of the registration job and the actions associated with the job.

      6 Use the following command to inspect the logs for the registration pod.

```
kubectl --namespace nsxi-platform logs --selector='job-name=cloud-connector-
register' --container=main
```

4 After resolving the error, click **Actions** in the NSX Network Detection and Response feature card. Select **Delete** to initiate the deletion of the partially activated NSX Network Detection and Response feature. After the delete process finishes, retry activating the feature again.

NSX Network Detection and Response Feature Activation Failed

The NSX Network Detection and Response feature activation wizard reports an error.

Problem

The NSX Network Detection and Response feature activation failed and the activation wizard reports an error similar to the following output.

`The feature activation took too long. Either the Kubernetes pods failed to come up or the registration with NSX Manager failed.`

The **Activate** button remains dimmed.

Cause

The NSX Network Detection and Response feature activation requires the deployment of several Kubernetes-based workloads on top of the NSX Application Platform cluster. If the cluster is in a degraded or an unstable state, the NSX Network Detection and Response activation can fail.

Solution

1 Ensure that NSX Application Platform is deployed correctly and is reported as `STABLE` on the **Systems > NSX Application Platform** UI page.

2 Inspect the logs for the NSX Cloud Connector registration service.

   a Collect an NSX Application Platform support bundle and inspect the logs for any Kubernetes pod with the name starting with `nsx-ndr-enable-ids` or `nsx-ndr-setup-kafka`.

   b Alternatively, the logs can also be queried interactively on the NSX Manager appliance using the following steps.

      1 Log into the NSX Manager appliance as **`root`**.

      2 Use the following command to mark the Kubernetes configuration for any subsequent `helm` and `kubectl` invocations.

```
export KUBECONFIG=/config/vmware/napps/.kube/config
```

      3 Using the following command, ensure that the NSX Cloud Connector helm chart is deployed successfully.

```
helm --namespace nsxi-platform list --all --filter 'nsx-ndr'
```

      Verify that the `STATUS` property displays `deployed`.

      4 Use the following command to inspect that the setup pods are deployed and completed successfully.

```
kubectl --namespace nsxi-platform get pods --selector='job-name in (nsx-ndr-enable-ids, nsx-ndr-setup-kafka)'
```

      Two pods must exist and both are showing the `STATUS` as `Completed`.

      5 Inspect the logs for the setup pod, using the following command.

```
kubectl --namespace nsxi-platform logs --selector='job-name in (nsx-ndr-enable-ids, nsx-ndr-setup-kafka)'
```

3 After resolving the error, click **Actions** in the NSX Network Detection and Response feature card. Select **Delete** to initiate the deletion of the partially activated NSX Network Detection and Response feature. After the delete process finishes, retry activating the feature again.

Time-Based Firewall Policy

With time windows, security administrators can restrict traffic from a source or to a destination, for a specific time period.

Time-based rules are available for distributed and gateway firewalls on ESXi hosts. Time windows apply to a firewall policy section, and all the rules in it. Each firewall policy section can have one time window. The same time window can be applied to more than one policy section. If you want the same rule applied on different days or different times for different sites, you must create more than one policy section. Time-based rules are available for distributed and gateway firewalls on ESXi hosts.

In NSX 4.0.1.1 and later, time-based rules are supported on both Local Managers and Global Managers in NSX Federation. Time can be specified in UTC for all sites, or time can be specified per local time zone. If you want the same rule applied on different days or different times for different sites, you must create more than one policy section.

Prerequisites

Network Time Protocol (NTP) is an Internet protocol used for clock synchronization between computer clients and servers. NTP service must be running on each transport node when using time-based rule publishing.

If a time-zone is changed on the edge transport node after the node is deployed, reload the edge node or restart the data plane for time-based gateway firewall policy to take effect. For details see Configuring NTP on Appliances and Transport Nodes.

Create a firewall policy.

Procedure

1 Click the clock icon on the firewall policy you want to have a time window.

A time window appears.

2 Click **Add New Time Window** and enter a **name**.

3 Select a time zone: UTC (Coordinated Universal Time), or the local time of the transport node. Distributed firewall only supports UTC with NTP service enabled, a change of time zone configuration is not supported.

4 Select the frequency of the time window - **Weekly** or **One time**.

5 Select the days of the week that the time window takes effect.

NSX supports configuring weekly UTC time-windows for the local time-zone, when the entire time-window for the local time-zone is within the same day as the UTC time-zone. For example, you cannot configure a time window in UTC for a 7am-7pm PDT, which maps to UTC 2pm-2am of the next day.

6 Select the beginning and ending dates for the time window, and the times the window will be in effect.

7 Click **Save.**

8 Click the check box next to the policy section you want to have a time window. Then click the clock icon.

9 Select the time window you want to apply, and click **Apply**.

10 Click **Publish**. The clock icon for the section turns green.

For the first publication of a time-based rule, the time is taken, and rule enforcement begins at less than 2 minutes. After the rules are deployed, enforcement as per time window, is instantaneous.

Troubleshooting Firewall

This section provides information about troubleshooting firewall issues.

Monitor and Troubleshoot Firewall on NSX Manager

There are several steps to take when troubleshooting firewall.

1 Check the Firewall policy realization status. See Check Rule Realization Status.

2 Check the rule hits statistics by navigating to **Security > Distributed Firewall** or **Security > Gateway Firewall**, and clicking the graph icon. Rule level stistics are aggregated every 15 minutes from all the transport nodes. Rule statistics can be reset using **Reset All Rules Stats** from the three dot menu icon .

3 Check for Capacity Dashboard to make sure configuration is within the supported limit of NSX. The Capacity dashboard can be accessed from **Security > Security Overview > Capacity** , see View the Usage and Capacity of Categories of Objects.

4 Check for supported configuration max limit for the given release by checking the Configuration Limits.

5 Check for per VM level Firewall Rules pushed to datapath in Manager Mode by navigating **Logical Switches > Ports > Related Firewall Rules**.

   You can also use the following NSX DFW helper script from github to get the total firewall rules configured and per VM firewall rules. https://github.com/vmware-samples/nsx-t/blob/master/helper-scripts/DFW/nsx-get-dfw-rules-per-vm.py

Troubleshooting Distributed Firewall on ESX Hosts

On ESX hosts, follow these steps to troubleshoot the NSX distributed firewall (DFW) data path issues.

Get the list of VMs on the ESXi host and associated Filter Name

This lists all VM's on this ESXi host. Note down the value of "name" field and use that in subsequent commands to get relevant output for a given VM.

```
[root@esxcomp-2a:~] summarize-dvfilter | grep -A 3 vmm
world 1371516 vmm0:PROD-MRS-DB-01 vcUuid:'50 20 92 e1 11 b7 10 d3-56 c5 e0 da 46 87 b5 d2'
 port 67108881 PROD-MRS-DB-01.eth0
  vNic slot 2
   name: nic-1371516-eth0-vmware-sfw.2
--
world 1622816 vmm0:DEV-MRS-DB-01 vcUuid:'50 2d f3 a3 96 a4 f4 94-6e 55 84 85 c1 bd 05 2c'
 port 67108883 DEV-MRS-DB-01.eth0
  vNic slot 2
   name: nic-1622816-eth0-vmware-sfw.2
--
world 7014985 vmm0:PROD-MRS-APP-01 vcUuid:'50 20 9b 5f cd b7 43 de-ab bb 8d 0e f5 bb ca 99'
 port 67108895 PROD-MRS-APP-01.eth0
  vNic slot 2
```

```
   name: nic-7014985-eth0-vmware-sfw.2
--
world 7022287 vmm0:PROD-MRS-APP-02 vcUuid:'50 20 4a 44 17 fb 21 cf-fb 62 1e a3 d0 3c 7d cf'
 port 67108896 PROD-MRS-APP-02.eth0
  vNic slot 2
   name: nic-7022287-eth0-vmware-sfw.2
[root@esxcomp-2a:~]
```

Get the firewall rules applied to a VM

Use Filter name associated with the VM from above output to get all the firewall rules applied to that VM's vNIC

```
[root@esxcomp-2a:~] vsipioctl getrules -f nic-7014985-eth0-vmware-sfw.2
ruleset mainrs {
  # generation number: 0
  # realization time : 2020-12-16T23:41:30
  # PRE_FILTER rules
  rule 5134 at 1 inout protocol any from addrset d8e7adac-af3b-4f22-9785-0cc30f0e81b1 to
addrset d8e7adac-af3b-4f22-9785-0cc30f0e81b1 accept with log tag 'ipv6-app-allow';
  rule 5133 at 2 inout protocol any from any to any accept with log tag 'ipv6-app-deny-
default';
  rule 5132 at 3 inout inet protocol icmp from any to addrset
9b14a216-4318-4bb1-94b0-56dfedec6f24 accept with log tag 'icmp-test';
  rule 5132 at 4 inout inet protocol tcp strict from any to addrset
9b14a216-4318-4bb1-94b0-56dfedec6f24 port 22 accept with log tag 'icmp-test';
  rule 5132 at 5 inout inet protocol ipv6-icmp from any to addrset
9b14a216-4318-4bb1-94b0-56dfedec6f24 accept with log tag 'icmp-test';
  rule 5130 at 6 inout inet protocol icmp from any to addrset rdst5130 accept with log tag
'icmp-test-gb-default';
  rule 5130 at 7 inout inet protocol ipv6-icmp from any to addrset rdst5130 accept with log
tag 'icmp-test-gb-default';
  # FILTER (APP Category) rules
  rule 5102 at 1 inout protocol any from addrset rsrc5102 to addrset d19f38e1-c13e-4fbb-9d6b-
b6971f251e2d accept;
  rule 5126 at 2 in protocol any from addrset rsrc5127 to addrset d19f38e1-c13e-4fbb-9d6b-
b6971f251e2d accept;
  rule 5127 at 3 out protocol any from addrset rsrc5127 to addrset d19f38e1-c13e-4fbb-9d6b-
b6971f251e2d accept;
  rule 5128 at 4 out protocol any from addrset rsrc5128 to addrset rdst5128 accept;
  rule 5129 at 5 in protocol any from addrset rsrc5128 to addrset
98abd76f-351b-4a4a-857f-1d91416b0798 accept;
  rule 5103 at 6 in protocol any from addrset rsrc5128 to addrset b1ed4d3d-ab4c-4bab-999b-
a50642cad495 accept;
  rule 5135 at 7 inout protocol any from any to any with attribute profile
acf76e7d-400b-438b-966f-8d5c10bebbda accept;
  rule 5135 at 8 inout protocol any from any to any with attribute profile 88dc6bf0-808e-49f6-
a692-dd0e5cee6ab3 accept;
  rule 5124 at 9 inout protocol any from any to any with attribute profile 8774c654-0f9e-43ad-
a803-4aa720e590cf accept;
  rule 5123 at 10 inout protocol any from any to any with attribute profile 13e599b5-
dd2d-420f-8473-9d45f0d324ac accept;
  rule 5125 at 11 inout protocol any from any to any with attribute profile e4be8d7e-
e4ab-4466-8f2e-998445ead95d accept;
```

```
  rule 2 at 12 inout protocol any from any to any drop with log tag 'icmp-default-rule';
}

ruleset mainrs_L2 {
  # generation number: 0
  # realization time : 2020-12-16T23:41:30
  # FILTER rules
  rule 1 at 1 inout ethertype any stateless from any to any accept;
}

[root@esxcomp-2a:~]
```

Get stats per FW rule per VM VNIC

Use "-s" with the above command to get the firewall stats associated with the VM firewall rules.

```
  [root@esxcomp-2a:~] vsipioctl getrules -f nic-7014985-eth0-vmware-sfw.2 -s
ruleset mainrs {
  # PRE_FILTER rules
rule  5134 at 1, 68 evals, 68 hits, 68 sessions, in 1120 out 1120 pkts, in 113952 out 114184
bytes
rule  5133 at 2, 24 evals, 24 hits, 24 sessions, in 16 out 8 pkts, in 896 out 768 bytes
rule  5132 at 3, 0 evals, 0 hits, 0 sessions, in 0 out 0 pkts, in 0 out 0 bytes
rule  5132 at 4, 0 evals, 0 hits, 0 sessions, in 0 out 0 pkts, in 0 out 0 bytes
rule  5132 at 5, 0 evals, 0 hits, 0 sessions, in 0 out 0 pkts, in 0 out 0 bytes
rule  5130 at 6, 0 evals, 0 hits, 0 sessions, in 0 out 0 pkts, in 0 out 0 bytes
rule  5130 at 7, 0 evals, 0 hits, 0 sessions, in 0 out 0 pkts, in 0 out 0 bytes
  # FILTER (APP Category) rules
rule  5102 at 1, 0 evals, 0 hits, 0 sessions, in 0 out 0 pkts, in 0 out 0 bytes
rule  5126 at 2, 0 evals, 0 hits, 0 sessions, in 0 out 0 pkts, in 0 out 0 bytes
rule  5127 at 3, 0 evals, 0 hits, 0 sessions, in 0 out 0 pkts, in 0 out 0 bytes
rule  5128 at 4, 0 evals, 0 hits, 0 sessions, in 0 out 0 pkts, in 0 out 0 bytes
rule  5129 at 5, 0 evals, 0 hits, 0 sessions, in 0 out 0 pkts, in 0 out 0 bytes
rule  5103 at 6, 0 evals, 0 hits, 0 sessions, in 0 out 0 pkts, in 0 out 0 bytes
rule  5135 at 7, 92 evals, 0 hits, 0 sessions, in 0 out 0 pkts, in 0 out 0 bytes
rule  5135 at 8, 92 evals, 0 hits, 0 sessions, in 0 out 0 pkts, in 0 out 0 bytes
rule  5124 at 9, 92 evals, 0 hits, 0 sessions, in 0 out 0 pkts, in 0 out 0 bytes
rule  5123 at 10, 92 evals, 0 hits, 0 sessions, in 0 out 0 pkts, in 0 out 0 bytes
rule  5125 at 11, 92 evals, 0 hits, 0 sessions, in 0 out 0 pkts, in 0 out 0 bytes
rule     2 at 12, 92 evals, 0 hits, 0 sessions, in 0 out 0 pkts, in 0 out 0 bytes
}

ruleset mainrs_L2 {
  # FILTER rules
rule     1 at 1, 0 evals, 0 hits, 0 sessions, in 0 out 0 pkts, in 0 out 0 bytes
}

[root@esxcomp-2a:~]
```

Get the addrset/groups used in the VM's Firewall rules

The firewall rule uses groups/addrset in the Source or destination. This output gets the all the addrset used in the rules based on the grouping configuration.

```
[root@esxcomp-2a:~] vsipioctl getaddrset -f nic-1371516-eth0-vmware-sfw.2
addrset is shared for this filter
global addrset
addrset 98abd76f-351b-4a4a-857f-1d91416b0798 {
ip 7.7.7.7,
ip 8.8.8.8,
}
addrset 9b14a216-4318-4bb1-94b0-56dfedec6f24 {
ip 10.1.0.0,
ip 10.2.0.2,
ip 10.114.217.26,
ip 172.16.202.2,
ip 172.16.202.22,
ip 192.168.202.2,
ip 2001::172:16:202:2,
ip 2001::172:16:202:22,
mac 00:50:56:a0:0e:25,
mac 00:50:56:a0:26:dc,
mac 00:50:56:a0:2d:c0,
mac 00:50:56:a0:8d:90,
}
addrset b1ed4d3d-ab4c-4bab-999b-a50642cad495 {
ip 7.7.7.7,
ip 8.8.8.8,
}
addrset d19f38e1-c13e-4fbb-9d6b-b6971f251e2d {
ip 3.3.3.3,
ip 4.4.4.4,
}
addrset d8e7adac-af3b-4f22-9785-0cc30f0e81b1 {
ip 172.16.202.2,
ip 172.16.202.22,
ip 2001::172:16:202:2,
ip 2001::172:16:202:22,
mac 00:50:56:a0:26:dc,
mac 00:50:56:a0:8d:90,
}
addrset rdst5128 {
ip 3.3.3.3,
ip 4.4.4.4,
ip 7.7.7.7,
ip 8.8.8.8,
}
addrset rdst5130 {
ip 1.1.1.1,
ip 1.1.1.2,
ip 100.100.100.100,
}
addrset rsrc5102 {
ip 1.1.1.1,
```

```
ip 1.1.1.2,
}
addrset rsrc5127 {
ip 1.1.1.1,
ip 1.1.1.2,
ip 3.3.3.3,
ip 4.4.4.4,
}
addrset rsrc5128 {
ip 1.1.1.1,
ip 1.1.1.2,
ip 3.3.3.3,
ip 4.4.4.4,
ip 7.7.7.7,
ip 8.8.8.8,
}
local addrset
No address sets.
[root@esxcomp-2a:~]
```

Get the active Firewall flow per VM

NSX DFW maintains active flow per VNIC. This output gets the all the active flows over that VNIC.

```
[root@esxcomp-2a:~] vsipioctl getflows -f nic-7014985-eth0-vmware-sfw.2
Count retrieved from kernel active=6, inactive=0, drop=0
ecbd448200000001 Active ipv6-icmp 86dd IN 5134   0 0  2001::172:16:202:22 ->
2001::172:16:202:2 128 0 1039376 1039376 9994 9994 tmo 9
ecbd4482000000b9 Active tcp 0800 OUT 5134*   0 0 (est) 172.16.202.2:Unknown(39914) ->
172.16.202.22:ssh(22) 305 EST:EST  rtt 21020 retrans 0/0 4409 3725 23 25 tmo 43195
ecbd4482000000ba Active ipv6-icmp 86dd OUT 5134*   0 0  fe80::250:56ff:fea0:8d90 ->
2001::172:16:202:22 135 0 64 72 1 1
ecbd4482000000bb Active igmp 0800 IN 5133*   0 0 (D) 0.0.0.0 -> 224.0.0.1 36 0 1 0 tmo 51
ecbd4482000000bc Active ipv6-icmp 86dd IN 5133*   0 0 (D) fe80::ffff:ffff:ffff:ffff ->
ff02::1 130 0 76 0 1 0 tmo 11
ecbd4482000000bd Active ipv6-icmp 86dd OUT 5133*   0 0 (D) fe80::250:56ff:fea0:8d90 ->
ff02::16 143 0 0 96 0 1 tmo 11
[root@esxcomp-2a:~]
```

Get the active Full Firewall config per VM

This output provides full firewall config per VNIC- Rules, Addrset & Profiles used.

```
[root@esxcomp-2a:~] vsipioctl getfwconfig -f nic-7014985-eth0-vmware-sfw.2
ruleset mainrs {
  # generation number: 0
  # realization time : 2020-12-16T23:41:30
  # PRE_FILTER rules
  rule 5134 at 1 inout protocol any from addrset d8e7adac-af3b-4f22-9785-0cc30f0e81b1 to
addrset d8e7adac-af3b-4f22-9785-0cc30f0e81b1 accept with log tag 'ipv6-app-allow';
  rule 5133 at 2 inout protocol any from any to any accept with log tag 'ipv6-app-deny-
default';
  rule 5132 at 3 inout inet protocol icmp from any to addrset
9b14a216-4318-4bb1-94b0-56dfedec6f24 accept with log tag 'icmp-test';
```

```
  rule 5132 at 4 inout inet protocol tcp strict from any to addrset
9b14a216-4318-4bb1-94b0-56dfedec6f24 port 22 accept with log tag 'icmp-test';
  rule 5132 at 5 inout inet protocol ipv6-icmp from any to addrset
9b14a216-4318-4bb1-94b0-56dfedec6f24 accept with log tag 'icmp-test';
  rule 5130 at 6 inout inet protocol icmp from any to addrset rdst5130 accept with log tag
'icmp-test-gb-default';
  rule 5130 at 7 inout inet protocol ipv6-icmp from any to addrset rdst5130 accept with log
tag 'icmp-test-gb-default';
  # FILTER (APP Category) rules
  rule 5102 at 1 inout protocol any from addrset rsrc5102 to addrset d19f38e1-c13e-4fbb-9d6b-
b6971f251e2d accept;
  rule 5126 at 2 in protocol any from addrset rsrc5127 to addrset d19f38e1-c13e-4fbb-9d6b-
b6971f251e2d accept;
  rule 5127 at 3 out protocol any from addrset rsrc5127 to addrset d19f38e1-c13e-4fbb-9d6b-
b6971f251e2d accept;
  rule 5128 at 4 out protocol any from addrset rsrc5128 to addrset rdst5128 accept;
  rule 5129 at 5 in protocol any from addrset rsrc5128 to addrset
98abd76f-351b-4a4a-857f-1d91416b0798 accept;
  rule 5103 at 6 in protocol any from addrset rsrc5128 to addrset b1ed4d3d-ab4c-4bab-999b-
a50642cad495 accept;
  rule 5135 at 7 inout protocol any from any to any with attribute profile
acf76e7d-400b-438b-966f-8d5c10bebbda accept;
  rule 5135 at 8 inout protocol any from any to any with attribute profile 88dc6bf0-808e-49f6-
a692-dd0e5cee6ab3 accept;
  rule 5124 at 9 inout protocol any from any to any with attribute profile 8774c654-0f9e-43ad-
a803-4aa720e590cf accept;
  rule 5123 at 10 inout protocol any from any to any with attribute profile 13e599b5-
dd2d-420f-8473-9d45f0d324ac accept;
  rule 5125 at 11 inout protocol any from any to any with attribute profile e4be8d7e-
e4ab-4466-8f2e-998445ead95d accept;
  rule 2 at 12 inout protocol any from any to any drop with log tag 'icmp-default-rule';
}

ruleset mainrs_L2 {
  # generation number: 0
  # realization time : 2020-12-16T23:41:30
  # FILTER rules
  rule 1 at 1 inout ethertype any stateless from any to any accept;
}

addrset is shared for this filter
global addrset
addrset 98abd76f-351b-4a4a-857f-1d91416b0798 {
ip 7.7.7.7,
ip 8.8.8.8,
}
addrset 9b14a216-4318-4bb1-94b0-56dfedec6f24 {
ip 10.1.0.0,
ip 10.2.0.2,
ip 10.114.217.26,
ip 172.16.202.2,
ip 172.16.202.22,
ip 192.168.202.2,
ip 2001::172:16:202:2,
ip 2001::172:16:202:22,
```

```
ip fe80::250:56ff:fea0:26dc,
ip fe80::250:56ff:fea0:8d90,
mac 00:50:56:a0:0e:25,
mac 00:50:56:a0:26:dc,
mac 00:50:56:a0:2d:c0,
mac 00:50:56:a0:8d:90,
}
addrset b1ed4d3d-ab4c-4bab-999b-a50642cad495 {
ip 7.7.7.7,
ip 8.8.8.8,
}
addrset d19f38e1-c13e-4fbb-9d6b-b6971f251e2d {
ip 3.3.3.3,
ip 4.4.4.4,
}
addrset d8e7adac-af3b-4f22-9785-0cc30f0e81b1 {
ip 172.16.202.2,
ip 172.16.202.22,
ip 2001::172:16:202:2,
ip 2001::172:16:202:22,
ip fe80::250:56ff:fea0:26dc,
ip fe80::250:56ff:fea0:8d90,
mac 00:50:56:a0:26:dc,
mac 00:50:56:a0:8d:90,
}
addrset rdst5128 {
ip 3.3.3.3,
ip 4.4.4.4,
ip 7.7.7.7,
ip 8.8.8.8,
}
addrset rdst5130 {
ip 1.1.1.1,
ip 1.1.1.2,
ip 100.100.100.100,
}
addrset rsrc5102 {
ip 1.1.1.1,
ip 1.1.1.2,
}
addrset rsrc5127 {
ip 1.1.1.1,
ip 1.1.1.2,
ip 3.3.3.3,
ip 4.4.4.4,
}
addrset rsrc5128 {
ip 1.1.1.1,
ip 1.1.1.2,
ip 3.3.3.3,
ip 4.4.4.4,
ip 7.7.7.7,
ip 8.8.8.8,
}
local addrset
```

```
No address sets.
containers are shared for this filter
global containers
container 13e599b5-dd2d-420f-8473-9d45f0d324ac {
# generation number: 21208
# realization time : 2020-12-16T23:41:30
FQDN : login\.microsoft\.com(3940c0d7-cbfc-abbb-35b4-786fc4199684),
}
container 8774c654-0f9e-43ad-a803-4aa720e590cf {
# generation number: 21208
# realization time : 2020-12-16T23:41:30
FQDN : outlook\.office365\.com(6e465c1d-7d81-9672-00e1-76ddfc280b8b),
}

container 88dc6bf0-808e-49f6-a692-dd0e5cee6ab3 {
# generation number: 21208
# realization time : 2020-12-16T23:41:30
APP_ID : APP_360ANTIV,
}

container acf76e7d-400b-438b-966f-8d5c10bebbda {
# generation number: 21208
# realization time : 2020-12-16T23:41:30
APP_ID : APP_ACTIVDIR,
}

container e4be8d7e-e4ab-4466-8f2e-998445ead95d {
# generation number: 21208
# realization time : 2020-12-16T23:41:30
FQDN : play\.google\.com(c44ef0fc-a922-eb1b-f155-4f0625271198),
}
local containers
No containers.
[root@esxcomp-2a:~]
```

Other output for FW troubleshooting

In addition to above command option NSX allows other options to debug the NSX FW datapath on ESX. Use the help menu as below.

```
[root@esxcomp-2a:~] vsipioctl -h
Usage: help <cmd> <options>
  below is a list of available cmd:
    getfilters      : get list of filters
    getfwconfig     : get rules, addrsets and containers of a filter
    getrules        : get rules of a filter
    getaddrsets     : get addrsets of a filter
    getcontainers   : get containers of a filter
    getspoofguard   : get spoofguard setting of a filter
    getflows        : get flows of a filter
    getconncount    : get active connection count
    getconnections  : get active connections
    getsisvmstats   : get service insertion service VM stats
    getsisvctable   : dump service insertion service table
```

```
    getsinshtable   : display service insertion nsh table
    getsiproxytable : display service insertion proxy table
    getsifailedspis : get service insertion failed spi table
    getsiflowprogtable : get service insertion flow programming table
    getsislotid     : get service insertion slot id
    getsilbenablestatus: get service insertion load balance enable status
    getmeminfo      : get meminfo data
    initvsiplogging : init vsip logger
    getfqdnentries  : get fqdn entries
    getdnsconfigprofile : get dns config profile for a filter
    getfilterstat   : get statistics of a filter
    gettimeout      : get connection timeout setting of a filter
    getfloodstat    : get flood protection status
    getsidcache     : get sid cache of a filter
    help            : this help message
  run `vsipioctl <cmd> -h' to find out available options of a cmd.
[root@esxcomp-2a:~]
```

NSX CLI for FW troubleshooting

On ESXi, nsxcli option can be used as an alternative option to ESX cli, by typing "nsxcli" and user can use "get firewall" command tree to get the similar output as above.

```
[root@esxcomp-2a:~] nsxcli
esxcomp-2a.dg.vsphere.local>
esxcomp-2a.dg.vsphere.local> get firewall
% Command not found: get firewall

  Possible alternatives:
    get firewall <vifuuid> addrsets
    get firewall <vifuuid> profile
    get firewall <vifuuid> ruleset rules
    get firewall exclusion
    get firewall ipfix-containers
    get firewall ipfix-filters
    get firewall ipfix-profiles
    get firewall ipfix-stats
    get firewall packetlog
    get firewall packetlog last <lines>
    get firewall rule-stats
    get firewall rule-stats total
    get firewall status
    get firewall thresholds
    get firewall vifs

esxcomp-2a.dg.vsphere.local> get firewall packetlog last 10
Wed Dec 16 2020 UTC 23:53:55.693
2020-12-16T23:53:23.878Z fd2e9266 INET6 match PASS 5134 OUT 72 ICMP fe80::250:56ff:fea0:8d90-
>fe80::250:56ff:fea0:26dc ipv6-app-allow
2020-12-16T23:53:23.878Z 5f46e9b1 INET6 match PASS 5134 IN 72 ICMP fe80::250:56ff:fea0:8d90-
>fe80::250:56ff:fea0:26dc ipv6-app-allow
2020-12-16T23:53:29.234Z fd2e9266 INET6 TERM 5134 OUT ICMP 135 0 fe80::250:56ff:fea0:8d90-
>2001::172:16:202:22 1/1 72/64 ipv6-app-allow
2020-12-16T23:53:29.234Z 5f46e9b1 INET6 TERM 5134 IN ICMP 135 0 fe80::250:56ff:fea0:8d90-
```

```
>2001::172:16:202:22 1/1 72/64 ipv6-app-allow
2020-12-16T23:53:30.234Z fd2e9266 INET6 TERM 5134 IN ICMP 135 0 fe80::250:56ff:fea0:26dc-
>fe80::250:56ff:fea0:8d90 1/1 72/64 ipv6-app-allow
2020-12-16T23:53:30.234Z 5f46e9b1 INET6 TERM 5134 OUT ICMP 135 0 fe80::250:56ff:fea0:26dc-
>fe80::250:56ff:fea0:8d90 1/1 72/64 ipv6-app-allow
2020-12-16T23:53:35.239Z fd2e9266 INET6 TERM 5134 OUT ICMP 135 0 fe80::250:56ff:fea0:8d90-
>fe80::250:56ff:fea0:26dc 1/1 72/64 ipv6-app-allow
2020-12-16T23:53:35.241Z 5f46e9b1 INET6 TERM 5134 IN ICMP 135 0 fe80::250:56ff:fea0:8d90-
>fe80::250:56ff:fea0:26dc 1/1 72/64 ipv6-app-allow
2020-12-16T23:53:51.876Z fd2e9266 INET6 match PASS 5134 OUT 72 ICMP fe80::250:56ff:fea0:8d90-
>2001::172:16:202:22 ipv6-app-allow
2020-12-16T23:53:51.876Z 5f46e9b1 INET6 match PASS 5134 IN 72 ICMP fe80::250:56ff:fea0:8d90-
>2001::172:16:202:22 ipv6-app-allow

esxcomp-2a.dg.vsphere.local> get firewall exclusion
Wed Dec 16 2020 UTC 23:53:57.731
                         Firewall Exclusion
------------------------------------------------------------------------

Exclusion count: 7
   00894e3c-8948-4b6b-a4cd-acd3a2c21205
   15f077e9-4492-4391-9f63-a99b6c978003
   2936443e-128c-4b6d-9fcf-3b2fad778b08
   3602f84a-8333-44f3-a3c2-e04fbf5e848f
   8149b7ec-553d-48e1-af04-1ee2f5ae266e
   d615679c-092e-4bfe-8c17-803fe8b3315d
   da619e9d-48a0-4c82-a831-bf580d3bec05

esxcomp-2a.dg.vsphere.local> get firewall thresholds
Wed Dec 16 2020 UTC 23:53:59.905
                              Firewall Threshold Monitors
------------------------------------------------------------------------------------------------

 #      Name      Raised  Threshold  CurrValue  CurrSize   MaxSize  PeakEver  EverTime(ago)
 1    dfw-cpu     False       60         0         --        --        0        --:--:--
 2   vsip-attr    False       60         3        4 MB     128 MB      3       4d 23:35:06
 3   vsip-flow    False       60         0        0 MB     312 MB      0        --:--:--
 4  vsip-fprules  False       60         0        0 MB     128 MB      0        --:--:--
 5   vsip-fqdn    False       60         0        0 MB     128 MB      0        --:--:--
 6  vsip-module   False       60         15      153 MB    1024 MB     15      4d 23:35:06
 7   vsip-rules   False       60         0        0 MB     512 MB      0        --:--:--
 8    vsip-si     False       60         0        0 MB     128 MB      0        --:--:--
 9   vsip-state   False       60         0        0 MB     384 MB      0        --:--:--

esxcomp-2a.dg.vsphere.local>
```

DFW L2 Rules Show Unknown MAC Address

After configuring a layer-2 firewall rule with one MAC set as source and another MAC set as destination, the getrules command on the host shows the destination MAC set as 01:00:00:00:00:00/01:00:00:00:00:00. For example,

```
[root@host1:~] vsipioctl getrules -f nic-1000052822-eth1-vmware-sfw.2
ruleset mainrs {
  # generation number: 0
  # realization time : 2018-07-26T12:42:28
  rule 1039 at 1 inout protocol tcp from any to any port 1521 accept as oracle;
  # internal # rule 1039 at 2 inout protocol tcp from any to any port 1521 accept;
  rule 1039 at 3 inout protocol icmp from any to any accept;
  rule 2 at 4 inout protocol any from any to any accept with log;
}

ruleset mainrs_L2 {
  # generation number: 0
  # realization time : 2018-07-26T12:42:28
  rule 1040 at 1 inout ethertype any stateless from addrset
d83a1523-0d07-4b18-8a5b-77a634540b57 to addrset 9ad9c6ef-c7dd-4682-833d-57097b415e41 accept;
  # internal # rule 1040 at 2 in ethertype any stateless from addrset
d83a1523-0d07-4b18-8a5b-77a634540b57 to addrset 9ad9c6ef-c7dd-4682-833d-57097b415e41 accept;
  # internal # rule 1040 at 3 out ethertype any stateless from addrset
d83a1523-0d07-4b18-8a5b-77a634540b57 to mac 01:00:00:00:00:00/01:00:00:00:00:00 accept;
  rule 1 at 4 inout ethertype any stateless from any to any accept;
}
```

The internal OUT rule with the address 01:00:00:00:00:00/01:00:00:00:00:00 is created by design to handle outbound broadcasting packets and does not indicate a problem. The firewall rule will work as configured.

Troubleshooting Gateway Firewall

Use the user interface and API to troubleshoot gateway firewall.

Use NSX Manager UI and API to check the following:

- Gateway Firewall is enabled for the given Gateway.
- Check the realization state for a given gateway firewall policy. The UI shows the realization status next to the top right side of the FW Policy header.
- Check rule stats to see any traffic is hitting the FW policy.
- Enable logging for the rule for troubleshooting the policy.

Gateway firewall is implemented on NSX Edge transport node. As a next step, use datapath troubleshooting as below using nsxcli commands on the NSX Edge node command prompt.

Get UUID of the Gateway on which Firewall is enabled

```
EDGE-VM-A01> get logical-router
Logical Router
UUID                                   VRF    LR-ID  Name
Type                        Ports
736a80e3-23f6-5a2d-81d6-bbefb2786666   0      0
TUNNEL                      4
8ccc0151-82bd-43d3-a2dd-6a31bf0cd29b   1      1      DR-DC-Tier-0-GW
DISTRIBUTED_ROUTER_TIER0    5
5a914d04-305f-402e-9d59-e443482c0e15   2      1025   SR-DC-Tier-0-GW
SERVICE_ROUTER_TIER0        7
495f69d7-c46e-4044-8b40-b053a86d157b   4      2050   SR-PROD-Tier-1
SERVICE_ROUTER_TIER1        5
```

Get all Gateway interfaces using UUID

Gateway firewall is implemented per Uplink interface of a Gateway. Identify the uplink interface and get the interface ID from the output below.

```
dc02-nsx-edgevm-1> get logical-router 16f04a64-ef71-4c03-bb5c-253a61752222 interfaces
Wed Dec 16 2020 PST 17:24:13.134
Logical Router
UUID                                   VRF    LR-ID  Name                                Type
16f04a64-ef71-4c03-bb5c-253a61752222   5      2059   SR-PROD-ZONE-GW
SERVICE_ROUTER_TIER1
Interfaces (IPv6 DAD Status A-DAD_Success, F-DAD_Duplicate, T-DAD_Tentative, U-
DAD_Unavailable)
    Interface     : 748d1f17-34d0-555e-8984-3ef9f9367a6c
    Ifuid         : 274
    Mode          : cpu
    Port-type     : cpu

    Interface     : 1bd7ef7f-4f3e-517a-adf0-846d7dff4e24
    Ifuid         : 275
    Mode          : blackhole
    Port-type     : blackhole

    Interface     : 2403a3a4-1bc8-4c9f-bfb0-c16c0b37680f
    Ifuid         : 300
    Mode          : loopback
    Port-type     : loopback
    IP/Mask       : 127.0.0.1/8;::1/128(NA)

    Interface     : 16cea0ab-c977-4ceb-b00f-3772436ad972         <<<<<<<<<< INTERFACE ID
    Ifuid         : 289
    Name          : DC-02-Tier0-A-DC-02-PROD-Tier-1-t1_lrp
    Fwd-mode      : IPV4_ONLY
    Mode          : lif
    Port-type     : uplink                                       <<<<<<<<<< Port-type Uplink
Interface
    IP/Mask       :
100.64.96.1/31;fe80::50:56ff:fe56:4455/64(NA);fc9f:aea3:1afb:d800::2/64(NA)
    MAC           : 02:50:56:56:44:55
```

```
    VNI           : 69633
    Access-VLAN   : untagged
    LS port       : be42fb2e-b10b-499e-a6a9-221da47a4bcc
    Urpf-mode     : NONE
    DAD-mode      : LOOSE
    RA-mode       : SLAAC_DNS_TRHOUGH_RA(M=0, O=0)
    Admin         : up
    Op_state      : up
    MTU           : 1500
    arp_proxy     :
```

Get Gateway Firewall Rules on a GW Interface

Use Interface ID to get firewall rules programmed on a gateway interface.

```
dc02-nsx-edgevm-2> get firewall 16cea0ab-c977-4ceb-b00f-3772436ad972 ruleset rules
Wed Dec 16 2020 PST 17:43:53.047
DNAT rule count: 0


SNAT rule count: 0


Firewall rule count: 6
    Rule ID   : 5137
    Rule      : inout protocol tcp from any to any port {22, 443} accept with log

    Rule ID   : 3113
    Rule      : inout protocol icmp from any to any accept with log

    Rule ID   : 3113
    Rule      : inout protocol ipv6-icmp from any to any accept with log

    Rule ID   : 5136
    Rule      : inout protocol any from any to any accept with log

    Rule ID   : 1002
    Rule      : inout protocol any from any to any accept

    Rule ID   : 1002
    Rule      : inout protocol any stateless from any to any accept

dc02-nsx-edgevm-2>
```

Check Gateway Firewall Sync status

Gateway Firewall sync flow status between Edge Nodes for high availability. Gateway firewall sync config can be seen using the output below.

```
dc02-nsx-edgevm-1> get firewall 16cea0ab-c977-4ceb-b00f-3772436ad972 sync config
Wed Dec 16 2020 PST 17:30:55.686
HA mode             : secondary-active
Firewall enabled    : true
Sync pending        : false
```

```
Bulk sync pending   : true        Last status: ok
Failover mode       : non-preemptive
Local VTEP IP       : 172.16.213.125
Peer VTEP IP        : 172.16.213.123
Local context       : 16f04a64-ef71-4c03-bb5c-253a61752222
Peer context        : 16f04a64-ef71-4c03-bb5c-253a61752222

dc02-nsx-edgevm-1>

dc02-nsx-edgevm-2> get firewall 16cea0ab-c977-4ceb-b00f-3772436ad972 sync config
Wed Dec 16 2020 PST 17:47:43.683
HA mode             : primary-passive
Firewall enabled    : true
Sync pending        : false
Bulk sync pending   : true        Last status: ok
Failover mode       : non-preemptive
Local VTEP IP       : 172.16.213.123
Peer VTEP IP        : 172.16.213.125
Local context       : 16f04a64-ef71-4c03-bb5c-253a61752222
Peer context        : 16f04a64-ef71-4c03-bb5c-253a61752222

dc02-nsx-edgevm-2>
```

Check Gateway Firewall Active Flows

Gateway firewall active flows can be seen using the command below. The flow states are synced between active and standby edge nodes for that gateway. The example below shows output from both edge-node-1 and edge-node-2.

```
dc02-nsx-edgevm-2> get firewall 16cea0ab-c977-4ceb-b00f-3772436ad972 connection
Wed Dec 16 2020 PST 17:45:55.889
Connection count: 2
0x0000000330000598: 10.166.130.107:57113 -> 10.114.217.26:22  dir in protocol tcp state
ESTABLISHED:ESTABLISHED fn 5137:0
0x04000003300058f1: 10.166.130.107 -> 10.114.217.26  dir in protocol icmp  fn 5136:0

dc02-nsx-edgevm-2>

dc02-nsx-edgevm-1> get firewall 16cea0ab-c977-4ceb-b00f-3772436ad972 connection
Wed Dec 16 2020 PST 17:47:09.980
Connection count: 2
0x0000000330000598: 10.166.130.107:57113 -> 10.114.217.26:22  dir in protocol tcp state
ESTABLISHED:ESTABLISHED fn 5137:0
0x04000003300058f1: 10.166.130.107 -> 10.114.217.26  dir in protocol icmp  fn 3113:0

dc02-nsx-edgevm-1>
```

Check Gateway Firewall Logs

Gateway firewall packet logs provide the gateway VRF and GW Interface information, along with flow details. Gateway firewall packet logs can be found in the file named firewallpkt.log in the /var/log directory.

Other Command Line Options for debugging Gateway Firewall

```
dc02-nsx-edgevm-2> get firewall 16cea0ab-c977-4ceb-b00f-3772436ad972

  Possible alternatives:
    get firewall <uuid> addrset name <string>
    get firewall <uuid> addrset sets
    get firewall <uuid> attrset name <string>
    get firewall <uuid> attrset sets
    get firewall <uuid> connection
    get firewall <uuid> connection count
    get firewall <uuid> connection raw
    get firewall <uuid> connection state
    get firewall <uuid> ike policy [<rule-id>]
    get firewall <uuid> interface stats
    get firewall <uuid> ruleset [type <rule-type>] rules [<ruleset-detail>]
    get firewall <uuid> ruleset [type <rule-type>] stats
    get firewall <uuid> sync config
    get firewall <uuid> sync stats
    get firewall <uuid> timeouts
    get firewall [logical-switch <uuid>] interfaces
    get firewall interfaces sync

dc02-nsx-edgevm-2>
```

Check Rule Realization Status

DFW rules can be created, updated, and deleted using both the UI and API.

Rule Realization Status on UI

You can see the rule realization status for DFW and Gateway firewall policies by navigating to **Security > Distributed Firewall** or **Security Gateway Firewall**, and checking the rule realization status reported by transport nodes.

There are four possible values for the rule realization status:

- Success
- Error
- In Progress
- Unknown

Rule Realization Status Through APIs

If the rule was created and enforced at relevant nodes, the realization status can be checked by following Policy Manager APIs.

To check realization status for all the entities created in policy manager run the command: `GET: https://<Policy Appliance IP>/policy/api/v1/infra/realized-state/realized-entities`The realized state of the object should be "REALIZED" and 'runtime_status' should be "SUCCESS"

For example, the query to check the realized state of <e2d4c010-96c8-11e9-8c0a-f7581ab92530> of security policy at the Policy manager level is <f96f27c0-92b8-11e9-96af-b5e746a259e7> is GET https://10.172.121.219/policy/api/v1/infra/realized-state/realized-entities?intent_path=/infra/domains/default/security-policies/f96f27c0-92b8-11e9-96af-b5e746a259e7/rules/e2d4c010-96c8-11e9-8c0a-f7581ab92530


```
{
"results": [
{
"extended_attributes": [],
"entity_type": "RealizedFirewallRule",
"intent_paths": [
"/infra/domains/default/security-policies/1-communication-560"
],
"resource_type": "GenericPolicyRealizedResource",
"id": "default.1-communication-560.3-communication-110",
"display_name": "default.1-communication-560.3-communication-110",
"description": "default.1-communication-560.3-communication-110",
"path": "/infra/realized-state/enforcement-points/default/firewalls/firewall-sections/
default.1-communication-560/firewall-rules/default.1-communication-560.3-communication-110",
"relative_path": "default.1-communication-560.3-communication-110",
"parent_path": "/infra/realized-state/enforcement-points/default/firewalls/firewall-sections/
default.1-communication-560",
"intent_reference": [],
"realization_specific_identifier": "1028",
"state": "REALIZED",
"alarms": [],
"runtime_status": "IN_PROGRESS",
"_create_user": "system",
"_create_time": 1561673625030,
"_last_modified_user": "system",
"_last_modified_time": 1561674044534,
"_system_owned": false,
"_protection": "NOT_PROTECTED",
"_revision": 6
}
],
"result_count": 1
}
```


To check the overall realized status of section of every rule in a section on the hypervisor run the command:GET https://<policy-mgr>/policy/api/v1/infra/realized-state/status?include_enforced_status=true&intent_path=<Security_policy_path>.

There are four possible values for the consolidated status:

- Success
- Error
- In Progress
- Unknown

Table 16-15. Consolidated Status

Transport Node 1 Overall Status	Transport Node 2 Overall Status	Consolidated Status
ERROR	ERROR	ERROR
ERROR	IN_PROGRESS	ERROR
ERROR	UNKNOWN	ERROR
IN_PROGRESS	IN_PROGRESS	IN_PROGRESS
IN_PROGRESS	UNKNOWN	IN_PROGRESS
SUCCESS	SUCCESS	SUCCESS
SUCCESS	ERROR	ERROR
SUCCESS	IN_PROGRESS	IN_PROGRESS
SUCCESS	UNKNOWN	UNKNOWN
UNKNOWN	UNKNOWN	UNKNOWN

Distributed Firewall Packet Logs

If logging is enabled for firewall rules, you can look at the firewall packet logs to troubleshoot issues.

The log file is `/var/log/dfwpktlogs.log` on ESXi hosts.

Table 16-16. Firewall Log File Variables

Variable	Possible Values
Filter hash	A number that can be used to get the filter name and other information.
AF Value	INET, INET6
Reason	■ match: Packet matches a rule. ■ bad-offset: Datapath internal error while getting packet. ■ fragment: The non-first fragments after they are assembled to the first fragment. ■ short: Packet too short (for example, not even complete to include an IP header, or TCP/UDP header). ■ normalize: Malformed packets that do not have a correct header or a payload. ■ memory: Datapath out of memory. ■ bad-timestamp: Incorrect TCP timestamp. ■ proto-cksum: Bad protocol checksum. ■ state-mismatch: TCP packets that do not pass the TCP state machine check. ■ state-insert: Duplicate connection is found. ■ state-limit: Reached the maximum number of states that a datapath can track. ■ SpoofGuard: Packet dropped by SpoofGuard. ■ TERM: A connection is terminated.

Table 16-16. Firewall Log File Variables (continued)

Variable	Possible Values
Action	■ PASS: Accept the packet. ■ DROP: Drop the packet. ■ NAT: SNAT rule. ■ NONAT: Matched the SNAT rule, but cannot translate the address. ■ RDR: DNAT rule. ■ NORDR: Matched the DNAT rule, but cannot translate the address. ■ PUNT: Send the packet to a service VM running on the same hypervisor of the current VM. ■ REDIRECT: Send the packet to network service running out of the hypervisor of the current VM. ■ COPY: Accept the packet and make a copy to a service VM running on the same hypervisor of the current VM. ■ REJECT: Reject the packet.
Rule set and rule ID	*rule set/rule ID*
Direction	IN, OUT
Packet length	*/length*
Protocol	TCP, UDP, ICMP, or PROTO (protocol number) For TCP connections, the actual reason that a connection is terminated is indicated after the keyword TCP. If TERM is the reason for a TCP session, then an extra explanation appears in the PROTO row. The possible reasons for terminating a TCP connection include: RST (TCP RST packet), FIN (TCP FIN packet), and TIMEOUT (idle for too long) In the example above, it is *RST*. So it means that there is a *RST* packet in the connection that must be reset. For non-TCP connections (UDP, ICMP or other protocols), the reason for terminating a connection is only TIMEOUT.
Source IP address and port	*IP address/port*
Destination IP address and port	*IP address/port*
TCP flags	S (SYN), SA (SYN-ACK), A (ACK), P (PUSH), U (URGENT), F (FIN), R (RESET)
Number of packets	Number of packets. 22/14 - in packets / out packets
Number of bytes	Number of bytes. 7684/1070 - in bytes/ out bytes

The following is a regular log sample for distributed firewall rules:

```
2018-07-03T19:44:09.749Z b6507827 INET match PASS mainrs/1024 IN 52 TCP 192.168.4.3/49627->192.168.4.4/49153 SEW

2018-07-03T19:46:02.338Z 7396c504 INET match DROP mainrs/1024 OUT 52 TCP 192.168.4.3/49676->192.168.4.4/135 SEW
```

```
2018-07-06T18:15:49.647Z 028cd586 INET match DROP mainrs/1027 IN 36 PROTO 2 0.0.0.0->224.0.0.1

2018-07-06T18:19:54.764Z 028cd586 INET6 match DROP mainrs/1027 OUT 143 UDP
fe80:0:0:0:68c2:8472:2364:9be/546->ff02:0:0:0:0:0:1:2/547
```

The elements of a DFW log file format include the following, separated by a space:

- timestamp:
- last eight digits of the VIF ID of the interface
- INET type (v4 or v6)
- reason (match)
- action (PASS, DROP, REJECT)
- rule set name/ rule ID
- packet direction (IN/OUT)
- packet size
- protocol (TCP, UDP, or PROTO #)
- SVM direction for netx rule hit
- source IP address/source port>destination IP address/destination port
- TCP flags (SEW)

For passed TCP packets there is a termination log when the session has ended:

```
2018-07-03T19:44:30.585Z 7396c504 INET TERM mainrs/1024 OUT TCP RST 192.168.4.3/49627-
>192.168.4.4/49153 20/16 1718/76308
```

The elements of a TCP termination log include the following, separated by a space:

- timestamp:
- last 8 digits of the VIF ID of the interface
- INET type (v4 or v6)
- action (TERM)
- ruleset name/ rule ID
- packet direction (IN/OUT)
- protocol (TCP, UDP, or PROTO #)
- TCP RST flag
- SVM direction for netx rule hit
- source IP address/source port>destination IP address/destination port
- IN packet count/OUT packet count (all accumulated)
- IN packet size/OUT packet size

The following is a sample of FQDN log file for distributed firewall rules:

```
2019-01-15T00:34:45.903Z 7c607b29 INET match PASS 1031 OUT 48 TCP 10.172.178.226/32808->23.72.199.234/80 S www.sway.com(034fe78d-5857-0680-81e4-d8da6b28d1b4)
```

The elements of an FQDN log include the following, separated by a space:

- timestamp:
- last eight digits of the VIF ID of the interface
- INET type (v4 or v6)
- reason (match)
- action (PASS, DROP, REJECT)
- ruleset name/ rule ID
- packet direction (IN/OUT)
- packet size
- protocol (TCP, UDP, or PROTO #) - for TCP connections, the actual reason that a connection is terminated is indicated after the following IP address
- source IP address/source port>destination IP address/destination port
- TCP flags - S (SYN), SA (SYN-ACK), A (ACK), P (PUSH), U (URGENT), F (FIN), R (RESET
- domain name/UUID where UUID is the binary internal representation of the domain name

The following is a sample of Layer 7 log file for distributed firewall rules:

```
2019-01-15T00:35:07.221Z 82f365ae INET match REJECT 1034 OUT 48 TCP 10.172.179.6/49818->23.214.173.202/80 S APP_HTTP

2019-01-15T00:34:46.486Z 7c607b29 INET match PASS 1030 OUT 48 UDP 10.172.178.226/42035->10.172.40.1/53 APP_DNS
```

The elements of a Layer 7 log include the following, separated by a space:

- timestamp:
- last eight digits of the VIF ID of the interface
- INET type (v4 or v6)
- reason (match)
- action (PASS, DROP, REJECT)
- ruleset name/ rule ID
- packet direction (IN/OUT)
- packet size

- protocol (TCP, UDP, or PROTO #) - for TCP connections, the actual reason that a connection is terminated is indicated after the following IP address
- source IP address/source port>destination IP address/destination port
- TCP flags - S (SYN), SA (SYN-ACK), A (ACK), P (PUSH), U (URGENT), F (FIN), R (RESET
- APP_XXX is the discovered application

Bare Metal Server Security

Secure workloads that are running on Windows Server 2016 bare metal servers.

Prerequisites

- Configure WinRM for Windows bare metal servers.
- Install Linux packages needed for Linux bare metal servers.

You can provide connectivity and security to applications or workloads between:

- Physical workloads (bare metal server) and virtual workloads
- Virtual workloads and physical workloads (bare metal server)
- Physical workloads (bare metal server) and physical workloads (bare metal server)

The workloads can be on overlay or VLAN-backed networks and the workloads must not be outside of the perimeter of a Windows Server 2016 bare metal server.

Before securing bare metal hosts, do the following:

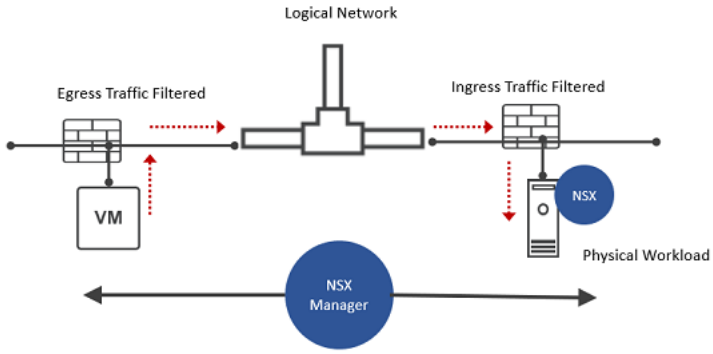
1 Ensure the NSX Agent is installed on the bare metal host.

2 Establish network connectivity between the application IP address of the Windows bare metal server, NSX Agent and NSX Manager.

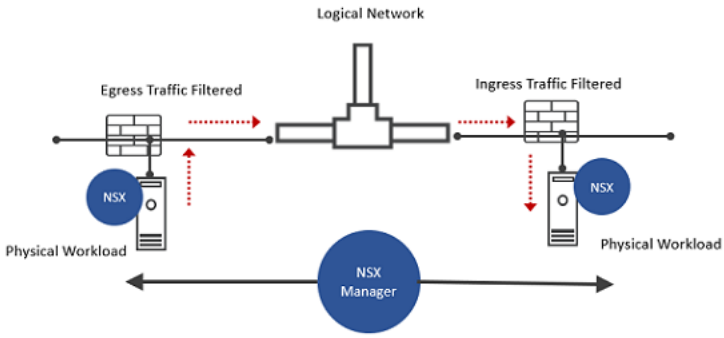
3 Apply DFW rules to secure ingress and egress traffic flowing through the L2 and L3 networks between workloads on a Windows Server 2016 bare metal server and virtual or physical workloads.

A couple of use cases where ingress and egress trafffic is filtered at the Windows bare metal server.

Traffic between Virtual and Physical Bare Metal Workloads



Traffic between Physical Bare Metal Workloads



Before you apply DFW rules to the Windows bare metal workloads, integrate NSX on the Windows Server using Ansible scripts. To install and integrate NSX on a Windows bare metal server, refer to the Secure Workloads on Windows Server 2016 Bare Metal Server topic in the *NSX Installation Guide*.

General Security Settings

Private IP Ranges

Private IP ranges are used to isolate suspicious traffic data within controlled network segments.

See Managing the Private IP Ranges for NSX Intelligence in *Using and Managing VMware NSX Intelligence*.

Firewall General Settings

This section contains profiles that fine tune Firewall Operations: Session Timers, Flood Protection, and DNS Security

Create a Session Timer

Session Timers define how long a session is maintained on the firewall after inactivity in the session.

When the session timeout for the protocol expires, the session closes. On the firewall, several timeouts for TCP, UDP, and ICMP sessions can be specified to apply to a user-defined group or a Tier-0 or Tier-1 gateway. Default session values can be modified depending on your network needs. Note that setting a value too low might cause frequent timeouts, and setting a value too high might delay failure detection. See Default Session Timer Values for more information.

Session timers are supported on ESXi hosts.

Procedure

1 Navigate to **Security > General Settings > Firewall > Session Timer**.

2 Click **Add Profile**.

The **Profile** screen appears, populated with the default values.

3 Enter a **name** and a **description** (optional) for the timer profile.

4 Click **Set** to select the Tier-0 or Tier-1 gateway or group to apply the timer profile.

5 Select the protocol. Accept the default values or enter your own values.

TCP Variables	Description
First Packet	The timeout value for the connection after the first packet has been sent. The default is 120 seconds.
Opening	The timeout value for the connection after a second packet has been transferred. The default is 30 seconds.
Established	The timeout value for the connection once the connection has become fully established.
Closing	The timeout value for the connection after the first FIN has been sent. The default is 120 seconds.
Fin Wait	The timeout value for the connection after both FINs have been exchanged and the connection is closed. The default is 45 seconds.
Closed	The timeout value for the connection after one endpoint sends an RST. The default is 20 seconds.

UDP Variables	Description
First Packet	The timeout value for the connection after the first packet is sent. This is the initial timeout for the new UDP flow. The default is 60 seconds.
Single	The timeout value for the connection if the source host sends more than one packet and the destination host has not sent one back. The default is 30 seconds.ESXi hosts only.
Multiple	The timeout value for the connection if both hosts have sent packets. The default is 60 seconds.

ICMP Variables	Description
First Packet	The timeout value for the connection after the first packet is sent. This is the initial timeout for the new ICMP flow. The default is 20 seconds.
Error reply	The timeout value for the connection after an ICMP error is returned in response to an ICMP packet. The default is 10 seconds. ESXi hosts only.

6 Click **Save**.

What to do next

After saving, click Manage Group to Profile Precedence to manage group to profile binding precedence.

Default Session Timer Values

The session timer profile applies the timeout values to Tier-0 or Tier-1 router interfaces or groups containing segments, segment-ports, tags, or any other non-IP based groups. The timeout values decide how long a protocol session remains active after the session closes.

Session Timer Values

- Default Timer Profile shown with API and UI applies only to distributed firewall (DFW).
- Gateway Firewall (GFW) default session timers are different than the default timer profile seen when using API and UI. GFW default session timers are optimized for North-South traffic, and some of them are lower than minimum configurable values by default.
- Firewall session timers can be changed for both DFW and GFW by using the API and UI.
- The same non-default timer profile can be applied to both DFW and GWF, if needed.

If you do not customize timer values, the gateway takes default values. Gateway firewall default timer values:

Timer Property	Edge Default (secs)	Minimum (secs)	Maximum (secs)
ICMP Error Reply	6	10	4320000
ICMP First Packet	6	10	4320000
TCP Closed	2	10	4320000
TCP Closing	900	10	4320000
TCP Established	7200	120	4320000
TCP Fin-wait	4	10	4320000
TCP First Packet	120	10	4320000
TCP Opening	30	10	4320000
UDP First Packet	30	10	4320000
UDP Multiple	30	10	4320000
UDP Single	30	10	4320000

Distributed firewall default session timer values:

Timer Property	DFW Default (secs)	Minimum (secs)	Maximum (secs)
ICMP Error Reply	10	10	4320000
ICMP First Packet	20	10	4320000

Timer Property	DFW Default (secs)	Minimum (secs)	Maximum (secs)
TCP Closed	20	10	4320000
TCP Closing	120	10	4320000
TCP Established	43200	120	4320000
TCP Fin-wait	45	10	4320000
TCP First Packet	120	10	4320000
TCP Opening	30	10	4320000
UDP First Packet	60	10	4320000
UDP Multiple	60	10	4320000
UDP Single	30	10	4320000

Flood Protection

Flood protection helps to protect against Denial of Service (DDoS) attacks.

DDoS attacks aim to make a server unavailable to legitimate traffic by consuming all the available server resources - the server is flooded with requests. Creating a flood protection profile imposes active session limits for ICMP, UDP, and half-open TCP flows. Distributed firewall can cache flow entries which are in SYN_SENT and SYN_RECEIVED states, and promote each entry to a TCP state after an ACK is received from the initiator, completing the three-way handshake.

Procedure

1 Navigate to **Security > General Settings > Firewall > Flood Protection**.

2 Click **Add Profile**, and select **Add Edge Gateway Profile** or **Add Firewall Profile**.

3 Fill out the flood protection profile parameters:

Table 16-17. Parameters for Firewall and Edge Gateway Profiles

Parameter	Minimum and maximum values	Default	
TCP Half Open Connection Limit - TCP SYN flood attacks are prevented by limiting the number of active, not-fully-established TCP flows which are allowed by the firewall.	1-1,000,000	Firewall - None Edge Gateway - 1,000,000	Set this text box to limit the number of active TCP half open connections. If this text box is empty, this limit is disabled on ESX nodes and set to the default on value of Edge Gateways.
UDP Active Flow Limit -UDP flood attacks are prevented by limiting the number of active UDP flows which are allowed by the firewall. Once the set UDP flow limit is reached, subsequent UDP packets which can establish a new flow are dropped.	1-1,000,000	Firewall - None Edge Gateway - 1,000,000	Set this text box to limit the number of active UDP connections. If this text box is empty, this limit is disabled on ESX nodes and set to the default on value of Edge Gateways.
ICMP Active Flow Limit - ICMP flood attacks are prevented by limiting the number of active ICMP flows which are allowed by the firewall. After the set flow limit is reached, subsequent ICMP packets which can establish a new flow are dropped.	1-1,000,000	Firewall - None Edge Gateway - 10,000	Set this text box to limit the number of active ICMP open connections. If this text box is empty, this limit is disabled on ESX nodes and set to the default on value of Edge Gateways.
Other Active Connection Limit	1-1,000,000	Firewall - None Edge Gateway - 10,000	Set this text box to limit the number of active connections other than ICMP, TCP, and UDP half open connections. If this text box is empty, this limit is disabled on ESX nodes, and set to the default on value of Edge Gateways.

Table 16-17. Parameters for Firewall and Edge Gateway Profiles (continued)

Parameter	Minimum and maximum values	Default	
SYN Cache - Syn Cache is used when a TCP half open connection limit has also been configured. The number of active half-open connections are enforced by maintaining a syncache of the not-fully-established TCP sessions. This cache maintains the flow entries which are in SYN_SENT and SYN_RECEIVED states. Each syncache entry will be promoted to a full TCP state entry after an ACK is received from the initiator, completing the three-way handshake.		Only available for firewall profiles.	Toggle on and off. Enabling SYN cache is effective only when a TCP half open connection limit is configured. Disabled by default.
RST Spoofing - Generates spoofed RST to server when purging half-open states from SYN cache. Allows server to clean up states associated with SYN flood (half open).		Only available for firewall profiles.	Toggle on and off. SYN Cache must be enabled for this option to be available
NAT Active Connection Limit	1 - 4294967295	Only available for Edge Gateway profiles. The default is 4294967295.	Set this parameter to limit the number of NAT connections that can be generated at the gateway.

4 To apply the profile to edge gateways and firewall groups, click **Set**.

5 Click **Save.**

What to do next

After saving, click Manage Group to Profile Precedence to manage group to profile binding precedence.

Configure DNS Security

Creating a DNS Security Profile helps to guard against DNS-related attacks.

You can do the following after you set up the DNS Security Profile:

- Snoop on DNS responses for a VM, or a group of VMs on the transport node to associate FQDN with IP addresses.

- Add global and default DNS server information, and apply it to all VMs that are using DFW rules.
- Specify selected DNS server information for selected VMs.
- Apply DNS profiles to groups.

Note Only ESXi is supported in the current release.

Procedure

1 Navigate to **Security > General Settings > Firewall > DNS Security**.

2 Click **Add Profile**.

3 Enter the following values:

Option	Description
Profile Name	Provide a profile name.
TTL	This field captures the Time to live for the DNS cache entry in seconds. You have the following options: TTL 0 - cached entry never expires. TTL 1 to 3599 - invalid TTL 3600 to 864000 – valid TTL left empty – automatic TTL, set from the DNS response packet. **Note** DNS Security Profile has a default DNS cache timeout of 24 hours.
Applied To	You can select a group based on any criteria to apply the DNS security profile to. **Note** Only one DNS server profile is applied to a VM.
Tags	Optional. Assign a tag and scope to the DNS profile to make it easy to search. See Add Tags to an Object for more information.

4 Click **Save**.

What to do next

After saving, click Manage Group to Profile Precedence to manage group to profile binding precedence.

Manage Group to Profile Precedence

You can bind multiple groups to a security profile. NSX applies the security profile to the group with highest precedence level.

If you bind a security profile to multiple groups, NSX assigns highest precedence to the newest group from that list. However, you can change the precedence level for groups.

To assign precedence to groups:

Prerequisites

- Session timer groups must only contain segments, segment ports, and VMs as members. Other category types are not supported.
- DNS security groups must contain only VMs as members. Other category types are not supported.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Navigate to **Security > General Settings > Firewall**.
3. Click **Manage Group to Profile Precedence**.
4. To assign a group highest level of precedence, move it to the top of the list.
5. Click **Close**.

Results

The security profile is applied to the group with highest precedence level.

Identity Firewall Event Log Sources

After configuring event log servers in the Active Directory, you need to turn on the Event Log Sources or Aria Operations for Logs.

When using event log scraping, ensure that NTP is correctly configured across all devices. See Time Synchronization between NSX Manager, vIDM, and Related Components.

Note Event log scraping enables IDFW for physical devices. Event log scraping can be used for virtual machines, however guest introspection will take precedence over event log scraping. Guest Introspection is enabled through VMware Tools and if you are using the complete VMware Tools installation and IDFW, guest introspection will take precedence over event log scraping.

Aria Operations for Logs 8.6 and later is supported with the provider configurations:

- Palo Alto Global Protect
- Aruba ClearPass

For more information about configuring Aria Operations for Logs (formerly known as vRealize Log Insight) see Integrate vRealize Log Insight with NSX Identity Firewall.

Navigate to **Security > General Settings > Identity Firewall Event Log Sources** and toggle the button for Event Log Sources or Aria Operations for Logs.

URL Database

URL database is used in TLS inspection, URL filtering, and FQDN analysis.

Enable the URL database by navigating to **Security > General Security Settings**, and selecting the URL Database tab.

Note that the latest URL data version (in the top left of the screen) may be different than the edge host's URL data version listed in the table.

Inventory

17

You can configure services, groups, context profiles, and virtual machines for the NSX inventory.

When you click the **Inventory** tab, an overview of the inventory objects is displayed, showing the number of groups, services, virtual machines, context profiles, L7 access profiles, and physical servers that are in the inventory. In addition, the following information about groups is shown:

- the number of groups used in policies
- the number of groups not used in policies
- the number of groups with members
- the number of groups without members
- the number of identity groups
- the number of identity groups used in policies
- the number of identity groups not used in policies

This chapter includes the following topics:

- Add a Service
- Add a Group
- Overview of Group Membership Criteria
- Profiles
- Attribute Types
- Containers
- Public Cloud Services
- Physical Servers
- Tags
- Resource Sharing
- Share Resources from the Default View

Add a Service

You can configure a service, and specify parameters for matching network traffic such as a port and protocol pairing.

You can also use a service to allow or block certain types of traffic in firewall rules. You cannot change the type after you create a service. Some services are predefined and cannot be modified or deleted.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Inventory > Services**.

3 Click **Add Service**.

4 Enter a name.

5 Click **Set Service Entries**.

6 Select a type.

The choices are **Layer 2** and **Layer 3 and above**.

7 Under **Port-Protocol**, click **Add Service Entry** to add one or more service entries.

For layer 2, the only available service type is **Ether**.

For layer 3 and above, the available service types are **IP**, **IGMP**, **ICMPv4**, **ICMPv6**,**ALG**, **TCP**, and **UDP**.

Note NSX supports the following built-in ALGs for DFW: FTP, TFTP, MS_RPC_TCP, MS_RPC_UDP, ORACLE_TNS, SUN_RPC_TCP and SUN_RPC_UDP.

NSX supports the following built-in ALGs for Gateway Firewall: FTP and TFTP.

8 Click the **Services** tab to add one or more services.

Any service that you add is considered a nested service because it is included in the service that you are creating. The recommended maximum level of nesting is 3. An example of three levels of nesting: service A includes service B, service B includes service C, and service C includes service D. In addition, cyclic nesting is not allowed. In the previous example, service C cannot include service A or B.

9 Click **Apply**.

10 (Optional) Add one or more tags.

11 (Optional) Enter a description.

12 Click **Save**.

Add a Group

Groups include different objects that are added both statically and dynamically, and can be used as the source and destination of a firewall rule.

Groups can be configured to contain a combination of virtual machines, IP sets, MAC sets, segment ports, segments, AD user groups, and other groups. Dynamic inclusion of groups can be based on tag, machine name, OS name, or computer name.

Note If you create a group in the API using LogicalPort based criteria, you cannot edit the group in the UI using the AND operator between SegmentPort criteria.

If Malicious IP Feed is enabled, a list of known malicious IPs are downloaded from NTICS cloud service. You can create groups to include these downloaded IPs and configure firewall rules to block access to them. Note that you can convert a Generic or IP Addresses Only groups to IP Addresses Only group with malicious IPs, but cannot convert the other way round.

Groups can also be excluded from firewall rules, and there are a maximum of 100 groups that can be on the list. IP sets, MAC sets, and AD groups cannot be included as members in a group that is used in a firewall exclusion list. See Manage a Firewall Exclusion List for more information.

A single group can be used as the source only within a distributed firewall rule. If IP and Active Directory groups are needed at the source, create two separate firewall rules.

Groups consisting of only IP addresses, MAC Addresses, or Active Directory groups cannot be used in the **Applied to** text box.

Note When a host is added to or removed from a vCenter Server, the external ID of the VMs on the host changes. If a VM is a static member of a group and the VM's external ID changes, the NSX Manager UI will no longer show the VM as a member of the group. However, the API that lists the groups will still show that the group contains the VM with its original external ID. If you add a VM as a static member of a group and the VM's external ID changes, you must add the VM again using its new external ID. You can also use dynamic membership criteria to avoid this issue.

For Policy Groups containing IPs, MAC addresses, and Identity Groups the listing API will **NOT** display the ‘members’ attribute. This applies to Groups containing a combination of static members also. For example, a Policy Group containing IP and VMs, will not display the the members attribute.

For Policy Groups not containing IPs, MAC addresses, or Identity Groups, the member attribute will be displayed in the NSGroup response. However new members and criteria introduced in NSX (such as DVPort and DVPG) will not be included in the MP group definition. Users can view the definition in Policy.

Tags in NSX are case-sensitive, but a group that is based on tags is "case- insensitive." For example, if the dynamic grouping membership criterion is `VM Tag Equals 'quarantine'`, the group includes all VMs that contain either the tags 'quarantine' or 'QUARANTINE'.

If you are using NSX Cloud, see Group VMs using NSX and Public Cloud Tags for information on the how to use public cloud tags to group your workload VMs in NSX Manager.

Prerequisites

If you are using NSX Federation, see Security in NSX Federation for details on configuration options.

Procedure

1. Select **Inventory > Groups** from the navigation panel.
2. Click **Add Group**, then enter a group name.
3. If you are adding a group from a Global Manager for NSX Federation, either accept the default region selection, or select a region from the drop-down menu. Once you create a group with a region, you cannot edit the region selection. However, you can change the span of the region itself by adding or removing locations from it. You can create customized regions before you create the group. See Create a Region from Global Manager.

   For groups added from a Global Manager in an NSX Federation environment, selecting a region is mandatory. This text box is not available if you are not using the Global Manager.
4. Click **Set**.

5 In the **Set Members** window, select the **Group Type**.

Group Type	Description
Generic	This group type is the default selection. A Generic group definition can consist of a combination of membership criteria, manually added members, IP addresses, MAC addresses, and Active Directory groups. Generic groups with only manually added IP address members are not supported for use in the **Applied To** field in DFW rules. It is possible to create the rule, but it will not be enforced. When you define membership criteria in the group, the members are dynamically added in the group based on one or more criteria. Manually added members include objects, such as segment ports, distributed ports, distributed port groups, VIFs, virtual machines, and so on.
IP Addresses Only	This group type contains only IP addresses (IPv4 or IPv6). **IP Addresses Only** groups with only manually added IP address members are not supported for use in the **Applied To** in DFW rules. It is possible to create the rule, but it will not be enforced. If the group type is **Generic**, you can edit its type to **IP Addresses Only** group or **IP Addresses Only** with malicious IPs group. In this case, only the IP addresses are retained in the group. All the membership criteria and other group definitions are lost. After a group of type **IP Addresses Only** or **IP Addresses Only** with malicious IPs is realized in NSX, you cannot edit the group type to **Generic**. **IP Addresses Only** group type is functionally similar to NSGroups with IP Set tag-based criterion in the Manager mode of earlier NSX releases.
IP Addresses Only with malicious IPs	If you have enabled Malicious IP Feeds, you can create an **IP Addresses Only** group with malicious IPs by switching on **Add Pre-Defined Malicious IPs**. For more information about configuring the feature, see Configure Malicious IP Feeds. You can also specify IPs and IP addresses only groups that should be treated as exceptions and must not be blocked. Note that once you have switched on the toggle **Add Pre-Defined Malicious IPs**, you cannot turn it off while editing the group.
Antrea	This group type is available only when your NSX environment has one or more Antrea Kubernetes clusters registered to it. For more information, see Antrea Groups and Add an Antrea Group.

6 (Optional) On the **Membership Criteria** page, click **Add Criterion** to add members in the generic group dynamically based on one or more membership criteria.

Starting in NSX 4.1, you can create generic groups with Kubernetes member types in dynamic membership criteria to match traffic entering into or leaving from registered Antrea Kubernetes clusters.

A membership criterion can have one or more conditions. The conditions can use the same member type or a mix of different member types. In a single membership criterion, conditions based on NSX member types cannot be mixed with conditions based on Kubernetes member types. However, you can have one criterion based on only NSX member types, and another criterion based on only Kubernetes member types, and then join both criteria with an OR operator.

Some restrictions apply to adding multiple conditions with mixed NSX member types or mixed Kubernetes member types in a membership criterion. To learn more about membership criteria, see Overview of Group Membership Criteria.

Note In a multi-tenant NSX environment, Kubernetes cluster resources are not exposed to the project inventory. Therefore, inside a project, you cannot create generic groups with Kubernetes member types in dynamic membership criteria.

7 (Optional) Click **Members** to add static members in the group.

The available member types are:

- **Groups** - If you are using NSX Federation, you can add a group as a member that has an equal or smaller span than the region you selected for the group you are creating from the Global Manager, see Security in NSX Federation
- **NSX Segments** - IP addresses assigned to a gateway interface, and NSX load balancer virtual IP addresses are not included as segment group members.
- **Segment Ports**
- **Distibuted Port Groups**
- **Distributed Ports**
- **VIFs**
- **Virtual Machines**
- **Physical Servers**
- **Cloud Native Service Instances**

8 (Optional) Click **IP/MAC Addresses** to add IP and MAC addresses as group members. IPv4 addresses, IPv6 addresses, and multicast addresses are supported.

Click **Action > Import** to import IP/MAC Addresses from a TXT file or a CSV file containing comma-separated IP/MAC values.

9 (Optional) Click **AD Groups** to add Active Directory Groups. Groups with Active Directory members can be used in the source text box of a distributed firewall rule for Identity Firewall. Groups can contain both AD and compute members.

Note If you are using NSX Federation, you cannot create groups from the Global Manager to include AD user groups.

10 (Optional) Enter a description and tag.

11 Click **Apply**

Groups are listed, with an option to view the members and where the group is used.

Overview of Group Membership Criteria

You can define membership criteria to add members dynamically in an NSX group based on one or more criteria.

A criterion can have one or more conditions. The conditions can use the same member type or a mix of different member types.

Some restrictions apply to adding multiple conditions with mixed NSX member types or mixed Kubernetes member types in a membership criterion. See the *Restrictions for Criteria with Mixed NSX Member Types* section and the *Restrictions for Criteria with Mixed Kubernetes Member Types* section later in this documentation.

By default, NSX uses the logical AND operator after each condition in a membership criterion. Other logical operators are not supported to join the conditions in a membership criterion.

To join criteria, OR and AND operators are available. By default, NSX selects the OR operator to join two criteria. AND operator is supported between two criteria only when:

- Both criteria use the same member type.
- Both criteria use a single condition.

The following restrictions apply to adding multiple conditions:

- A maximum of five conditions with the same member type is supported in a single membership criterion. For example, in a criterion, you can add a maximum of five conditions with the Virtual Machine member type.
- A maximum of 15 conditions with mixed member types are supported in a single membership criterion. For example, in a criterion, you can add a maximum of 15 conditions with a mix of NSX Segment and Segment Port member types.
- A maximum of 35 conditions with mixed member types are supported in a generic group.

A group can have a maximum of five membership criteria. However, the total number of criteria that you can add in a group is determined by the number of conditions in each criterion. See the following examples.

Example 1

A group with three membership criteria and a total of 35 conditions:

- Criterion 1 has 15 conditions with mixed member types.
- Criterion 2 has 15 conditions with mixed member types.
- Criterion 3 has 5 conditions with the same member type.

Example 2

A group with four membership criteria and a total of 35 conditions:

- Criterion 1 has 15 conditions with mixed member types.
- Criterion 2 has 14 conditions with mixed member types.
- Criterion 3 has four conditions with the same member type.
- Criterion 4 has two conditions with the same member type.

Example 3

A group with five membership criteria and a total of 22 conditions:

- Criterion 1 has 10 conditions with mixed member types.
- Criterion 2 has three conditions with the same member type.
- Criterion 3 has four conditions with the same member type.
- Criterion 4 has three conditions with the same member type.
- Criterion 5 has two conditions with mixed member types.

Because this group has reached the limit of five criteria, you cannot add another membership criterion. However, you can add more conditions, if required, in any of the five criteria until you don't exceed the following upper limits mentioned earlier:

- A maximum of five conditions with the same member type in a single criterion.
- A maximum of 15 conditions with mixed member types in a single criterion.
- A total of 35 conditions in the generic group.

Restrictions for Criteria with Mixed NSX Member Types

Member Type	Criterion With Mixed Member Types	Tag Operator	Scope Operator
Virtual Machine	Not Supported	■ Equals - one tag can be selected. ■ Contains ■ Starts with ■ Ends with	■ Equals
NSX Segment	Supported Conditions based on NSX Segment can be mixed with conditions based on Segment Port	■ Equals - one tag can be selected. ■ Not Equals - one tag can be selected.	■ Equals ■ Not Equals - if selected, the tag operator is removed.
Segment Port	Supported Conditions based on Segment Port can be mixed with conditions based on NSX Segment	■ Equals - one tag can be selected. ■ Not Equals - one tag can be selected. ■ Not In - a maximum of five tags can be selected.	■ Equals ■ Not Equals - if selected, the tag operator is removed.

Member Type	Criterion With Mixed Member Types	Tag Operator	Scope Operator
Distributed Port Groups	Supported Conditions based on Distributed Port Group can be mixed with conditions based on Distributed Port	■ Equals - one tag can be selected. ■ Not Equals - one tag can be selected.	■ Equals ■ Not Equals - if selected, the tag operator is removed.
Distributed Ports	Supported Conditions based on Distributed Port can be mixed with conditions based on Distributed Port Group	■ Equals - one tag can be selected. ■ Not Equals - one tag can be selected. ■ Not In - a maximum of five tags can be selected.	■ Equals ■ Not Equals - if selected, the tag operator is removed.
IP Set - This member type will be deprecated in the future. It is currently available to achieve backward compatibility with preexisting NSGroups or Groups based on IP Set tag-based criterion. We recommend you to use **Group** as the member type and add tag-based groups of type 'IP Addresses Only' in a membership criterion.	Not Supported	■ Equals - one tag can be selected.	■ Equals
Group - Use this member type to add tag-based groups of type 'IP Addresses Only' in a membership criterion.	Not Supported	Equals - one tag can be selected.	Equals

Overview of Membership Criteria with Kubernetes Member Types

Starting in NSX 4.1, you can create generic groups with Kubernetes member types in dynamic membership criteria to match traffic entering into or leaving from Antrea Kubernetes clusters.

You can then use these generic groups in distributed firewall rules or gateway firewall rules to secure traffic between VMs in the NSX environment and pods in Antrea Kubernetes clusters.

In a generic group, Kubernetes member types are available on the **Membership Criteria** page only when at least one Antrea Kubernetes cluster is registered to your NSX environment.

Note In a multi-tenant NSX environment, Kubernetes cluster resources are not exposed to the project inventory. Therefore, inside a project, you cannot create generic groups with Kubernetes member types in dynamic membership criteria.

The following table lists the Kubernetes member types that are available in the **Default** view of NSX 4.1 onwards to add dynamic membership criteria in a generic group. Each Kubernetes member type belongs to either cluster scope or namespace scope.

Member Type	Scope
Kubernetes Cluster	Cluster
Kubernetes Namespace	Namespace
Kubernetes Node	Cluster
Kubernetes Service	Namespace
Kubernetes Ingress	Namespace
Kubernetes Gateway	Namespace
Antrea Egress	Cluster
Antrea IP Pool	Cluster

Naming Conventions Used for Conditions With Kubernetes Member Types

The following table explains the naming conventions that are used in this documentation to represent the different conditions that you can add in dynamic membership criteria, which are based on Kubernetes member types.

Naming Convention for Condition	Meaning
Kubernetes Cluster condition	Conditions in the dynamic membership criteria are based on the Kubernetes Cluster member type.
Kubernetes Namespace condition	Conditions in the dynamic membership criteria are based on the Kubernetes Namespace member type.
Resource condition	Conditions in the dynamic membership criterion are based on any of these Kubernetes member types: ■ Kubernetes Service ■ Kubernetes Ingress ■ Kubernetes Gateway ■ Antrea Egress ■ Antrea IP Pool ■ Kubernetes Node
Cluster-scoped resource condition	Conditions in the dynamic membership criteria are based on any of these Kubernetes member types: ■ Antrea Egress ■ Antrea IP Pool ■ Kubernetes Node
Namespace-scoped resource condition	Conditions in the dynamic membership criteria are based on any of these Kubernetes member types: ■ Kubernetes Service ■ Kubernetes Ingress ■ Kubernetes Gateway

Restrictions for Criteria with Mixed Kubernetes Member Types

The following table provides a high-level summary of the restrictions or validations that apply to using Kubernetes member types in a single membership criterion. For examples of validations, see the *Validations for Dynamic Grouping with Kubernetes Member Types* section later in this documentation.

Member Type	Restrictions for Using the Member Type in a Criterion	Supported Attributes	Tag Operator	Scope Operator
Kubernetes Cluster	Cannot be used alone in a criterion. Only one cluster condition is allowed in a criterion. Must be mixed with at least one Kubernetes resource condition. Optionally, can be mixed with Kubernetes Namespace conditions and Kubernetes resource conditions.	Name	Not Supported	Not Supported
Kubernetes Namespace	Cannot be used alone in a criterion. Cannot be mixed with cluster-scoped resource conditions. Must be mixed with namespace-scoped resource conditions. Optionally, can be mixed with a Kubernetes Cluster condition.	Name Tag	Equals - one tag can be selected	Equals
Antrea Egress	Can be used alone in a criterion. Optionally, can be mixed with a Kubernetes Cluster condition.	Name Tag	Equals - one tag can be selected	Equals
Antrea IP Pool	Can be used alone in a criterion. Optionally, can be mixed with a Kubernetes Cluster condition.	Name Tag	Equals - one tag can be selected	Equals
Kubernetes Ingress	Can be used alone in a criterion. Optionally, can be mixed with only Kubernetes Namespace conditions or only Kubernetes Cluster condition, or both.	Name Tag	Equals - one tag can be selected	Equals
Kubernetes Gateway	Can be used alone in a criterion. Optionally, can be mixed with only Kubernetes Namespace conditions or only Kubernetes Cluster condition, or both.	Name Tag	Equals - one tag can be selected	Equals

Member Type	Restrictions for Using the Member Type in a Criterion	Supported Attributes	Tag Operator	Scope Operator
Kubernetes Service	Can be used alone in a criterion. Optionally, can be mixed with only Kubernetes Namespace conditions or only Kubernetes Cluster condition, or both.	Name Tag	Equals - one tag can be selected	Equals
Kubernetes Node	Can be used alone in a criterion. Only a single node condition is allowed. Optionally, can be mixed with a Kubernetes Cluster condition. Cannot be mixed with Kubernetes Namespace conditions or Kubernetes resource conditions.	IP Address Pod CIDR	Not supported See Note after this table.	Not supported

Note When you create a generic group with Kubernetes Node member type by using the API, only the Matches operator is allowed. This operator can take only the * value. The * wildcard character will match all the nodes in the K8s cluster (if the Kubernetes Node condition is mixed with a Kubernetes Cluster condition), or it will match nodes across all clusters (if the Kubernetes Node condition is used alone).

Validations for Dynamic Grouping with Kubernetes Member Types

Validation 1

A membership criterion can have a maximum of one Kubernetes Cluster condition. To match a single Kubernetes cluster by name, use the Equals operator and enter the cluster name.

Note Kubernetes cluster name must be unique.

If you want to match multiple clusters in a single Kubernetes Cluster condition, you can use one of these operators:

- In
- Starts With
- Ends With

Example:

Matches Single K8s Cluster	Matches Multiple K8s Clusters
Criterion: Kubernetes Cluster Name Equals ClusterA	Criterion: Kubernetes Cluster Name In ClusterA,ClusterB,ClusterC A maximum of five comma-separated values are allowed. The values must not be separated with spaces.

Validation 2

A membership criterion with a Kubernetes Cluster condition can be mixed with any one of the Kubernetes resource conditions. If you add Kubernetes Namespace conditions too in the same criterion, then resource conditions must be limited only to namespace-scoped resource conditions.

Examples:

- A membership criterion with a Kubernetes Cluster condition and three Kubernetes Ingress conditions is valid.
- A membership criterion with a Kubernetes Cluster condition and two Antrea Egress conditions is valid.
- A membership criterion with a Kubernetes Cluster condition and three Antrea IP Pool conditions is valid.
- A membership criterion with a Kubernetes Cluster condition, one Kubernetes Namespace condition, and three Kubernetes Gateway conditions is valid.

Validation 3

A membership criterion must include at least one Kubernetes resource condition. A membership criterion is invalid if it contains:

- Only Kubernetes Cluster condition
- Only Kubernetes Namespace condition
- Only Kubernetes Cluster condition and Kubernetes Namespace condition

Examples:

- A membership criterion with one Kubernetes Cluster condition, one Kubernetes Namespace condition, and one Kubernetes Ingress condition is valid.
- A membership criterion with one Kubernetes Cluster condition, two Kubernetes Namespace conditions, and three Kubernetes Ingress conditions is valid.

Validation 4

A membership criterion can have only a single Kubernetes Node condition. Optionally, a Kubernetes Node condition can be mixed with only a Kubernetes Cluster condition.

A Kubernetes Node condition cannot be mixed with Kubernetes Namespace conditions or Kubernetes resource conditions.

A membership criterion with only a Kubernetes Node condition can stand alone. However, a group with only a Kubernetes Node condition will match nodes of all Antrea Kubernetes clusters that are registered to NSX.

Tag operator and Scope operator are currently not supported for the Kubernetes Node condition.

Kubernetes Node condition supports the following two properties.

Property	Description
IP Address	Matches the internal IP address of all nodes of specified Antrea Kubernetes clusters.
Pod CIDR	Matches the Pod CIDR of all nodes of specified Antrea Kubernetes clusters.

Validation 5

If a membership criterion includes Kubernetes Cluster condition and Kubernetes Namespace conditions, then it must include at least one namespace-scoped Kubernetes resource condition. You cannot mix any of the following cluster-scoped Kubernetes resource conditions in the same criterion:

- Antrea Egress
- Antrea IP Pool
- Kubernetes Node

Examples:

- A membership criterion with one Kubernetes Cluster condition, two Kubernetes Namespace conditions, and Kubernetes Gateway conditions is valid.
- A membership criterion with one Kubernetes Cluster condition, four Kubernetes Namespace conditions, and three Kubernetes Service conditions is valid.
- A membership criterion with one Kubernetes Cluster condition, one Kubernetes Namespace condition, and one Antrea Egress condition is invalid because Antrea Egress is a cluster-scoped resource.

Validation 6

A membership criterion must include at least one Kubernetes resource condition. A resource condition can stand alone in a criterion. However, if you add multiple resource conditions in a criterion, then all resource conditions must be of the same member type.

Note Kubernetes Cluster condition and Kubernetes Namespace conditions are used for defining the scope of a criterion. They are not Kubernetes resource conditions, and hence they are not limited by this validation rule.

Examples:

- A membership criterion with five Kubernetes Service conditions is valid.
- A membership criterion with one Kubernetes Cluster condition, three Kubernetes Namespace conditions, and four Kubernetes Service conditions is valid.
- A membership criterion with one Kubernetes Cluster condition, three Kubernetes Namespace conditions, four Kubernetes Service conditions, and three Kubernetes Ingress conditions is invalid. The reason is that you mixed resource conditions of two different member types (Kubernetes Service and Kubernetes Ingress) in the same criterion.

However, you can create separate criteria with a resource condition based on a single member type and then join both criteria with an OR operator, as shown below:

Criterion 1:

One Kubernetes Cluster condition + three Kubernetes Namespace conditions + four Kubernetes Service conditions

OR

Criterion 2:

One Kubernetes Cluster condition + three Kubernetes Namespace conditions + three Kubernetes Ingress conditions

Validation 7

In a single membership criterion, conditions based on NSX member types cannot be mixed with conditions based on Kubernetes member types. However, you can have a group with one criterion based on only NSX members types, and other criterion based on only Kubernetes member types, and then join both criteria with an OR operator.

Example:

Valid	Invalid
Criterion 1: Virtual Machine conditions OR Criterion 2: Kubernetes Cluster condition + Kubernetes Gateway conditions	Criterion: NSX Segment condition + Segment Port condition AND Kubernetes Cluster condition + Kubernetes Gateway conditions

Profiles

There are two types of profiles: context profiles and layer 7 access profiles.

Profiles enable creating attributes key value pairs such as layer 7 App ID, and domain names. After a profile has been defined, it can be used in one or more distributed firewall rules, or gateway firewall rules.

Context Profiles

Context Profiles are used in firewall rules.

There are five attributes for use in context profiles: App ID, Custom URL, Domain (FQDN) Name, URL Category, and URL Reputation. Select App IDs can have one or more sub attributes, such as TLS_Version and CIPHER_SUITE. Both App ID and FQDN can be used in a single context profile. Multiple App IDs can be used in the same profile. One App ID with sub attributes can be used - sub attributes are cleared when multiple App ID attributes are used in a single profile.

Both system defined and user defined Fully Qualified Domain Names (FQDNs) are supported. You can see the list of FQDNs when you add a new context profile of attribute type *Domain (FQDN) Name*. You can also see a list of FQDNs, and where they are used by navigating to **Inventory > Context Profiles > FQDNs**.

Procedure

1 Select **Inventory > Profiles**.

2 Select the **Context Profile** tab and click **Add Context Profile**.

3 Enter a **Profile Name**, and optional **Description**.

4 In the Attributes column, click **Set**.

5 Click **Add Attribute**, and select one or more attributes from the drop-down menu: **App ID**, **Custom URL**, **Domain (FQDN) Name**, **URL Category**, or **URL Reputation**.

Attribute	Procedure
App ID - Advanced App IDs found here NSX Application IDs require a custom profile.	a Enter the name of the advanced App ID you'd like to use in firewall rules. b Click **Add**. c Click **Apply**.
Custom URL	To create a custom URL: a Click the three dot menu and select **Add Custom URL**. b Enter the URL. c Click **Save**.
Domain (FQDN) Name	Select a system FQDN by scrolling down the list. To create a user-defined FQDN: a Click the three dot menu and select **Add FQDN**. b Enter the domain name in the form **`*.[hostname].[domain]`**. For example, *.abracadabra.com. Do not include http:// or any other header. c Click **Save**. The newly created FQDN appears in the attribute value column. d Search and add additional FQDNs. e Click **Apply**.
URL Category	a Select one or more URL categories by scrolling down the list. b Click **Add**. c Click **Apply**.
URL Reputation	a Select one or more of the attributes by clicking in the box. b Click **Add**. c Click **Apply**. See FQDN Analysis Dashboard for more information about URL reputation.

6 (Optional) If you have selected an attribute with sub attributes such as SSL or CIFS, click **Set** in the Sub Attributes/Values column.

   a Click **Add Sub Attribute** and select TLS_VERSION, TLS_CIPHER_SUITE, or CIFS_SMB_VERSION.

   b Select one or more sub attributes.

   c Click **Add**. Another sub attribute can be added by clicking **Add Sub Attribute**.

   d Click **Apply**.

7 (Optional) Enter a tag or scope. See Tags for more information.

8 Click **Save**.

What to do next

Apply this context profile to a layer 7 distributed firewall rule (for layer 7 or domain name), or gateway firewall rule (for layer 7).

L7 Access Profiles

An L7 Access Profile can contain multiple entries with different attribute types: APP ID, URL category, custom URL, and URL reputation. This is used for URL filtering in a Gateway Firewall rule, in the profile field.

Layer 7 access profile has a default entry, and a user can add more entries above the default entry. Entries in the profile are evaluated in the listing order, and action is taken upon the first match.

Procedure

1 Select **Inventory > Profiles**.

2 Select the **L7 Access Profiles** tab and click **Add L7 Access Profile**.

3 Enter a **Profile Name**, and optional **Description**.

4 In the Attributes column, click **Set**.

5 Click **Add Attribute Type**, and select one or more attributes from the drop-down menu:

Attribute Type	Attribute Value
App ID - over 750	To view available App IDs, scroll down the list, or select App IDs.
URL Category - 80+ categories including social media, banking, phishing, etc.	Select one or more URL categories by scrolling down the list.
Custom URL - with regular expression	For more details see Custom URLs.
URL Reputation	Choose one or more of these reputations: High Risk, Suspicious, Moderate Risk, Low Risk, Trustworthy, Unknown

6 Select the rule action.

- Allow - allows matched traffic.
- Reject - rejects matched traffic.
- Reject with Response - rejects and sends the client the response page. This option is not available for the App ID attribute type. Navigate to **Security > Gateway Firewall > Settings > URL Filtering** to view and customize the **Reject with Response** message.

  The **Reject with Response** page is sent only for http traffic. The response page will contain the URL (first 10 bytes show), Category, Source IP and message-text. Enter the message for the **Reject with Response** page. Click **Preview Page** to view the page that will be sent when access to a URL is blocked by a policy.

7 By default, logging is turned off. Toggle the button to activate logging.

8 By default the entry is turned on. Toggle the button to deactivate the entry.

9 Click **Add**.

10 Click **Apply.**

Attribute Types

App IDs

Layer 7 attributes (App IDs) identify which application a particular packet or flow is generated by, independent of the port that is being used. Using App IDs reduces north south and east west attacks by only allowing appropriate traffic across an open port.

Enforcement based on App IDs enable users to allow or deny applications to run on any port, or to force applications to run on their standard port. vDPI enables matching packet payload against defined patterns, commonly referred to as signatures. Signature-based identification and enforcement enables customers to match the particular application/protocol a flow belongs to, and the version of that protocol, for example TLS version 1.0, TLS version 1.2 or different versions of CIFS traffic. This allows you to have visibility into or restrict the use of protocols that have known vulnerabilities for all deployed applications, and their E-W flows within the datacenter.

Layer 7 App IDs are used in context profiles and L7 access profiles in distributed firewall and gateway firewall rules.

Note NFS version 4 is not a supported attribute.

- Gateway firewall rules do not support the use of FQDN attributes or other sub attributes in context profiles.
- Context profiles are not supported on tier-0 gateway firewall policy.

Supported App IDs and FQDNs:

- For FQDN, users need to configure a high priority rule with a DNS App ID for the specified DNS servers on port 53.
- SYSLOG App ID is detected only on standard ports.

Below is a table with the list of Basic App IDs. For Advanced App IDs see NSX Application IDs.

Attribute (App ID)	Description	Type
360ANTIV	360 Safeguard is a program developed by Qihoo 360, an IT company based in China	Web Services
ACTIVDIR	Microsoft Active Directory	Networking
AMQP	Advanced Messaging Queuing Protocol is application layer protocol which supports business message communication between applications or organizations	Networking
AVAST	Traffic generated by browsing Avast.com official website of Avast! Antivirus downloads	Web Services
AVG	AVG Antivirus/Security software download and updates	File Transfer
AVIRA	Avira Antivirus/Security software download and updates	File Transfer
BLAST	A remote access protocol that compresses, encrypts, and encodes a computing experiences at a data center and transmits it across any standard IP network for VMware Horizon desktops.	Remote Access
BDEFNDER	BitDefender Antivirus/Security software download and updates.	File Transfer
CA_CERT	Certification authority (CA) issues digital certificates which certifies the ownership of a public key for message encryption	Networking
CIFS	CIFS (Common Internet File System) is used to provide shared access to directories, files, printers, serial ports, and miscellaneous communications between nodes on a network	File Transfer
CLDAP	Connectionless Lightweight Directory Access Protocol is an application protocol for accessing and maintaining distributed directory information services over an Internet Protocol (IP) network using UDP.	Networking
CTRXCGP	Citrix Common Gateway Protocol is an application protocol for accessing and maintaining distributed directory information services over an Internet Protocol (IP) network using UDP.	Database
CTRXGOTO	Hosting Citrix GoToMeeting, or similar sessions based on the GoToMeeting platform. Includes voice, video, and limited crowd management functions	Collaboration
CTRXICA	ICA (Independent Computing Architecture) is a proprietary protocol for an application server system, designed by Citrix Systems	Remote Access
DCERPC	Distributed Computing Environment / Remote Procedure Calls, is the remote procedure call system developed for the Distributed Computing Environment (DCE)	Networking

Attribute (App ID)	Description	Type
DIAMETER	An authentication, authorization, and accounting protocol for computer networks	Networking
DHCP	Dynamic Host Configuration Protocol is a protocol used management for the distribution of IP addresses within a network	Networking
DNS	Querying a DNS server over TCP or UDP	Networking
EPIC	Epic EMR is an electronic medical records application that provides patient care and healthcare information.	Client Server
ESET	Eset Antivirus/Security software download and updates	File Transfer
FPROT	F-Prot Antivirus/Security software download and updates	File Transfer
FTP	FTP (File Transfer Protocol) is used to transfer files from a file server to a local machine	File Transfer
GITHUB	Web-based Git or version control repository and Internet hosting service	Collaboration
HTTP	(HyperText Transfer Protocol) the principal transport protocol for the World Wide Web	Web Services
HTTP2	Traffic generated by browsing websites that support the HTTP 2.0 protocol	Web Services
IMAP	IMAP (Internet Message Access Protocol) is an Internet standard protocol for accessing email on a remote server	Mail
KASPRSKY	Kaspersky Antivirus/Security software download and updates	File Transfer
KERBEROS	Kerberos is a network authentication protocol designed to provide strong authentication for client/server applications by using secret-key cryptography	Networking
LDAP	LDAP (Lightweight Directory Access Protocol) is a protocol for reading and editing directories over an IP network	Database
MAXDB	SQL connections and queries made to a MaxDB SQL server	Database
MCAFEE	McAfee Antivirus/Security software download and updates	File Transfer
MSSQL	Microsoft SQL Server is a relational database.	Database
NFS	Allows a user on a client computer to access files over a network in a manner similar to how local storage is accessed. **Note** NFS version 4 is not a supported attribute.	File Transfer
NNTP	An Internet application protocol used for transporting Usenet news articles (netnews) between news servers, and for reading and posting articles by end user client applications.	File Transfer
NTBIOSNS	NetBIOS Name Service. In order to start sessions or distribute datagrams, an application must register its NetBIOS name using the name service	Networking

Attribute (App ID)	Description	Type
NTP	NTP (Network Time Protocol) is used for synchronizing the clocks of computer systems over the network	Networking
OCSP	An OCSP Responder verifying that a user's private key has not been compromised or revoked	Networking
ORACLE	An object-relational database management system (ORDBMS) produced and marketed by Oracle Corporation.	Database
PANDA	Panda Security Antivirus/Security software download and updates.	File Transfer
PCOIP	A remote access protocol that compresses, encrypts, and encodes a computing experiences at a data center and transmits it across any standard IP network.	Remote Access
POP3	Microsoft's implementation of NetBIOS Name Service (NBNS), a name server and service for NetBIOS computer names.	Mail
RADIUS	Provides centralized Authentication, Authorization, and Accounting (AAA) management for computers to connect and use a network service	Networking
RDP	RDP (Remote Desktop Protocol) provides users with a graphical interface to another computer	Remote Access
RTCP	RTCP (Real-Time Transport Control Protocol) is a sister protocol of the Real-time Transport Protocol (RTP). RTCP provides out-of-band control information for an RTP flow.	Streaming Media
RTP	RTP (Real-Time Transport Protocol) is primarily used to deliver real-time audio and video	Streaming Media
RTSP	RTSP (Real Time Streaming Protocol) is used for establishing and controlling media sessions between end points	Streaming Media
SIP	SIP (Session Initiation Protocol) is a common control protocol for setting up and controlling voice and video calls	Streaming Media
SMTP	SMTP (Simple Mail Transfer Protocol) An Internet standard for electronic mail (e-mail) transmission across Internet Protocol (IP) networks.	Mail
SNMP	SNMP (Simple Network Management Protocol) is an Internet-standard protocol for managing devices on IP networks.	Network Monitoring
SSH	SSH (Secure Shell) is a network protocol that allows data to be exchanged using a secure channel between two networked devices.	Remote Access
SSL	SSL (Secure Sockets Layer) is a cryptographic protocol that provides security over the Internet.	Web Services
SYMUPDAT	Symantec LiveUpdate traffic, this includes spyware definitions, firewall rules, antivirus signature files, and software updates.	File Transfer
SYSLOG	SYSLOG is a protocol that allows network devices to send event messages to a logging server.	Network Monitoring

Attribute (App ID)	Description	Type
TELNET	A network protocol used on the Internet or local area networks to provide a bidirectional interactive text-oriented communications facility using a virtual terminal connection.	Remote Access
TFTP	TFTP (Trivial File Transfer Protocol) being used to list, download, and upload files to a TFTP server like SolarWinds TFTP Server, using a client like WinAgents TFTP client.	File Transfer
VNC	Traffic for Virtual Network Computing.	Remote Access
WINS	Microsoft's implementation of NetBIOS Name Service (NBNS), a name server and service for NetBIOS computer names.	Networking

FQDNs

A fully qualified domain name (FQDN) is the complete domain name for a specific host on the Internet. FQDNs are used in firewall rules to allow or reject traffic going to specific domains.

The FQDN attribute type is used in distributed firewall FQDN Filtering policy, see FQDN Filtering . NSX supports custom FQDNs that are defined by an administrator in addition to the pre-defined list of FQDNs.

Note Custom FQDNs do not support custom top level domain names.

Custom FQDN supports the following:

- Starting in 4.0.1, FQDN supports processing of DNS response record packets containing canonical names (CNAMEs).
- Full FQDN names such as maps.google.com or myapp.corp.com
- Partial REGEX with * at the beginning only such as *eng.northpole.com or *yahoo.com
- FQDN name length up to 64 characters
- FQDN names must end with the registered top level domain (TLD) such as .com, .org, or .net

When creating a custom FQDN, using a wildcard domain is best practice. For example, using **`*.example.com`**, would include sub domains such as `americas.example.com` and `emea.example.com`. Using `example.com`, would not include any sub domains.

Procedure

1. From your browser, log in with admin privileges to an NSX Manager at https://<nsx-manager-ip-address>.
2. Select **Inventory > Profiles**.
3. Select the **Attribute Types** tab, and **FQDNs**.

   A table of system-generated FQDNs appears.
4. Select **Actions > Add FQDN**.

5 Enter the domain name in form **`*[hostname].[domain]`**. For example, *abracadabra.com

   Do not include http:// or any other header.

6 Click **Save**.

   The user-defined FQDN is shown in the table of available FQDNs, with User in the **Created By** column.

7 (Optional) To display a subset of FQDNs, click **Filter by Name, Path and more** and select **Created by** or **Domain**.

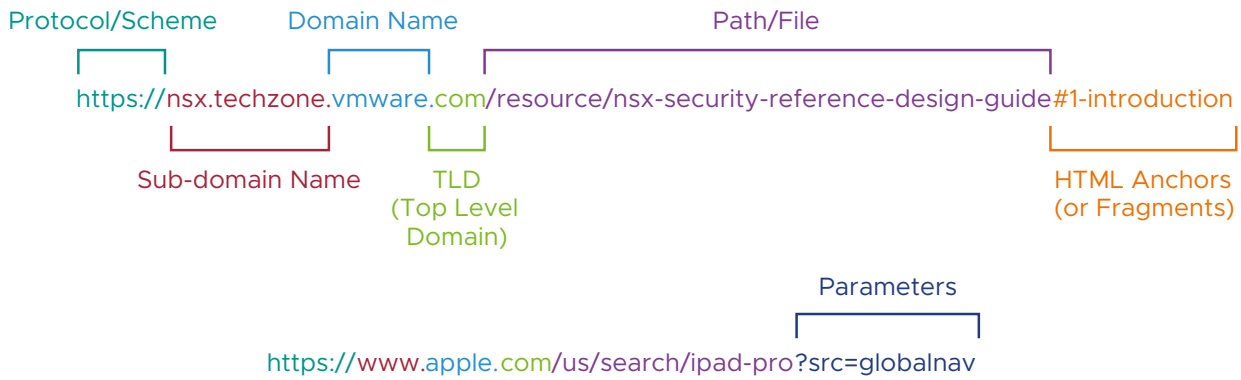
What to do next

FQDNs can be used in context profiles for distributed firewall rules.

Custom URLs

Custom URLs are used in L7 access profiles as an attribute type for URL filtering, and context profiles.

NSX allows users to configure full or partial domain or URL with special characters. The following image depicts different elements of a URL (Uniform Resource Locator).


Protocol/Scheme
Domain Name
Path/File
https://nsx.techzone.vmware.com/resource/nsx-security-reference-design-guide#1-introduction
Sub-domain Name
TLD
(Top Level
Domain)
HTML Anchors
(or Fragments)
Parameters
https://www.apple.com/us/search/ipad-pro?src=globalnav

NSX supports the following custom URL patterns:

- Only IANA registered Top Level Domain (TLD), and wildcard characters * and ^ are supported.
- Entered URL should be a valid URL. For example, www...google+com will be rejected.
- Do not include http://, https://, ftp://, etc.
- Users must enter a domain name, but may optionally include URI as well. For example, google.com and google.com/news are valid, but /news is invalid input.
- URI must not include query parameters. For example, www./google.com/query?keyboard=news will be rejected.
- Special characters * and ^ can be used for wildcard matching both in the domain name and the URI path. * will match one or more words where as ^ will match exactly one word. For example, *.google.com will match news.google.com and also local.news.google.com. However, ^ will only match news.google.com, but will not match local.news.google.com. Similarly, google.com/* will match google.com/news and also google.com/news/sanjose. However, google.com/^ will only match google.com/news, but will not match google.com/news/sanjose.
- * and ^ match only full words, they cannot be used to match partial words. For example, *google.com cannot be used to match mygoogle.com. However, *.google.com can be used to match my.google.com. Similarly, google.com/*news and google.com/n* are invalid and will be rejected.
- * and ^ can appear more than once in the URL. For example, *.google.* and *.google.com/^/news/* are valid.
- As ^ matches only a single word, two ^s back to back is allowed. For example, ^.^.google.com is valid and matches local.news.google.com, but does not match my.local.news.google.com or news.google.com. Similarly, google.com/^/^ is valid and matches google.com/news/sanjose, but does not match google.com/news or google.com/news/sanjose/today.
- When entering just the domain name (without URI), user may enter the domain name with a / at the end or without it and they have different behavior. If the user enters without an ending /, then it is treated as a partial match. For example, *.google.com will match news.google.com, news.google.com.us, news.google.com.us/, news.google.com.us/local, etc. If the user enters the domain name with a / at the end, then it is treated as an exact match. For example, *.google.com/ will match news.google.com and news.google.com/, but will not match news.google.com.us or news.google.com/local.
- When the entered URL has a path, then the presence of / at the end is not treated in any special way and the URL is treated as an exact match. For example, google.com/news/ will only match google.com/news/, but will not match google.com/news or google.com/newslatest or google.com/news/latest.

- URL matching can be used to match HTTP URL (Host header + URI) or TLS SNI. If the user specified URL has a path, it will not match the TLS SNI. However, if TLS Inspection is enabled and the traffic is decrypted, then the internal HTTP URL can be used for matching URL with path. For example, *.google.com and *.google.com/ can match HTTP URL or TLS SNI (without TLS Inspection), but *.google.com/news will not match TLS SNI (without TLS Inspection).

Examples

User Input	Example
espn.com	Matches: espn.com, espn.com/, espn.com.us, espn.com.us/, espn.com/sports, espn.com.us/sports/p Does not match: premium.espn.com
espn.com/	Matches: espn.com, espn.com/ Does not match: premium.espn.com, espn.com.us, espn.com.us/, espn.com/sports, espn.com.us/sports/
espn.com/sports	Matches: espn.com/sports Does not match: espn.com/sportsnba, esp.com/sports/, espn.com/sports/nba
espn.com/sports/	Matches: espn.com/sports Does not match: esp.com/sports, espn.com/sports/nba
*espn.com	Matches: premium.espn.com, replay.preimum.espn.com, instant.replay.premium.espn.com, premium.espn.com/, premium.espn.com.us, premium.espn.com.us/, premium.espn.com/sports, premium.espn.com.us/sports Does not match: espn.com, .espn.com
.espn.	Matches: www.espn.com, premium.espn.com, www.espn.us, premium.espn.com.us/, latest.espn.com/sports, www.espn.com.us/ sports/ replay.premium.espn.com, instant.replay.premium.espn.com instant.replay.premium.espn.com.us/ Does not match: espn.com, www.espn
espn.*.us/	Matches: espn.news.us, espn.news.us/, espn.local.news.us Does not match: espn.us, www.espn.us, espn.news.us/ sports
.espn./*	Matches: www.espn.com/sports, replay.premium.espn.com.us/sports/nba/ Does not match: www.espn.com, www.espn.com/
^.espn.com	Matches: www.espn.com, www.espn.com/, www.espn.com/sports, www.espn.com.us/ Does not match: espn.com, news.local.espn.com
^.espn.^	Matches: www.espn.com, www.espn.com/,www.espn.com.us, www.espn.com.us/ sports Does not match: news.local.espn.com, www.espn.com.us

User Input	Example
espn.^/	Matches: espn.com, espn.com/ Does not match: espn.com.us, espn.com/sports
www.^.^.com/^/^	Matches www.local.espn.com/sports/nba Does not match www.personal.local.espn.com/sports/nba, www.local.espn.com/sports/nba/

URL Categories

URL categories are used by in context profiles and L7 access profiles.

To view available URL categories and where they are used, navigate to **Inventory > Profiles > Attribute Types** and select the **URL Categories** tab. Click **Where Used** to view the attribute entries that contain the URL category.

Containers

You can view container objects in the NSX inventory and do basic diagnostic or troubleshooting tasks on container clusters in the NSX Manager UI. High level details about container objects such as, Namespaces, Pods, Kubernetes Services, Network Policies, and so on, are displayed in the UI.

If the **Containers** page is active in your browser, and some changes happened in the container cluster inventory in the background, the inventory details in the UI are not auto-refreshed. You must manually click the **Refresh** icon at the bottom of the page a few times to ensure that inventory updates are pushed to the NSX Manager UI.

View Details of Namespaces

You can filter the list of namespaces in the table using several criteria, such as ID, Label, Cluster Name, CNI Type, and many other criteria. At a high level, for each namespace, the table shows the number of Kubernetes Services, Pods, the name of container cluster in which the namespace is created, CNI type of container cluster, and so on.

Type of container cluster includes: Kubernetes, OpenShift, AKS, EKS, GKE, TKGm, to name a few. CNI Type includes: NSX Container Plugin (NCP), Antrea.

For details about Kubernetes Services and Pods in a specific namespace, click the hyperlinked number in the respective columns.

For example, if you want to view the list of all Pods that are associated with a specific Kubernetes Service, click the number in the **Pods** column. The **Pods Details** window opens. This pop-up window shows the status of Pods, Pod IP address, Labels associated with Pod, and other details. For container clusters with NCP as the CNI, only IP addresses of the Pod's corresponding segment ports are shown. If the network interfaces of the Pods are not attached to NSX segments, the Pod IP addresses are not shown.

Expand a row in the table to view more details about the namespace. For example:

- Number of Ingress Rules and the details of each Ingress Rule.

- Number of Network Policies.
- Number of Labels and the details of each Label.

This list does not mention all the inventory details that you can view for a namespace. Check the **Namespaces** page in the NSX Manager UI to know more.

Note

- For Antrea container clusters, the **Namespaces** page shows information about Kubernetes Network Policies and Antrea Network Policies, which have their scope limited to a namespace.
- The **Networking** and **Networking Status** columns show information only for container clusters that use NCP as the CNI. For container clusters with Antrea as the CNI, these two columns are not applicable.

View Details of Container Clusters

You can filter the list of container clusters in the table using several criteria, such as External ID, Cluster Name, CNI Type, and many other criteria. At a high level, for each container cluster, the table shows the number of Kubernetes Services, Pods, and Nodes, CNI Type, and Networking information.

For details about Kubernetes Services, Pods, Nodes, and Networking in a specific container cluster, click the hyperlinked number in the respective columns. The **Networking Status** column shows information only for container clusters that use NCP as the CNI. For container clusters with Antrea as the CNI, this column is not applicable.

Expand a row in the table to view more details about the container cluster. For example:

- Infrastructure type of the container cluster (example values: vSphere, AWS, Azure, Google, VMC, and so on).
- Antrea version (for Antrea container clusters).
- Number of Network Policies in the container cluster (for Antrea container clusters, this number refers to Antrea Cluster Network Policies).
- Number of namespaces in the cluster and the details of each namespace.

This list does not mention all the inventory details that you can view for a container cluster. Check the **Clusters** page in the NSX Manager UI to know more.

Related Documentation

To learn about NSX Container Plugin, go to NCP documentation at https://docs.vmware.com/en/VMware-NSX-T-Data-Center/index.html.

To learn about integrating Antrea container clusters to NSX, go to Chapter 24 Integration of Kubernetes Clusters with Antrea CNI.

Public Cloud Services

You can see a list of public cloud services that are available for your public cloud workload VMs.

Public cloud services that can be protected using cloud native security constructs, can be on-boarded with NSX Cloud.

Note In the current release, only the following AWS services are supported:

- RDS
- Application ELB (network ELB not supported)

How to onboard public cloud services using NSX Cloud

You onboard public cloud services in the same way as you onboard workload VMs in the Native Cloud Enforced Mode.

Once onboarded, the public cloud services are available from **Inventory > Public Cloud Services**.

You can create firewall rules for these services in the same way as for workload VMs.

See Managing VMs in the Native Cloud Enforced Mode.

Note You must enable the ports used by these services, for example, port 80 for ELB, in the firewall rules in NSX Manager.

Physical Servers

You can view the inventory of physical servers by navigating to **Inventory > Physical Servers**. These are transport nodes running on bare-metal servers.

For each physical server, the following information is displayed:

- Name
- OS type
- IP address
- Tags

Tags

Tags help you to label NSX objects so that you can quickly search or filter objects, troubleshoot and trace, and do other related tasks.

You can create tags using both the UI and APIs. Each tag has the following two attributes:

- Tag (refers to the tag name. It is required, must be unique and case-sensitive.)
- Scope (optional)

Tag scope is analogous to a key and tag name is analogous to a value. For example, let us say, you want to label all virtual machines based on their operating system (Windows, Mac, Linux). You can create three tags, such as Windows, Linux, and Mac, and set the scope of each tag to OS. Other examples of tag scope can be tenant, owner, name, and so on.

After you save a tag, you cannot update the name and scope. However, you can unassign or remove tags from objects.

For information about the maximum number of tags supported in NSX objects, see the VMware Configuration Maximums tool at https://configmax.vmware.com/home.

Following are some of the operations that you can do with tags:

- Assign or unassign tags to an object.
- Assign or unassign a single tag to multiple objects simultaneously (supported only for VMs).
- View a list of all tags in the inventory.
- Filter the list of tags by tag name, tag source, and tag scope.
- View a list of objects that are assigned a specific tag.

Note When a VM disappears from the VMware vCenter inventory for more than 30 minutes, NSX tags on the VM are lost. If the same VM reappears in the VMware vCenter inventory after 30 minutes, NSX treats it as a new VM, and you must add the tags again on the VM. This behavior is expected and as per design. For example, this behavior is seen when using array-based replication with VMware Site Recovery Manager™.

Use Cases of Tags

The following table describes some use cases of using tags.

Use Case	Description
Manageability	■ Simplify searching of objects in a large-scale inventory management. ■ Provide more information to differentiate objects that share similar or unclear names.
Third-party sharing and context sharing	■ Annotate objects with custom information. ■ Allow third-party non-NSX systems to add metadata information in an automated fashion. For example, metadata from partners, cloud management providers, container platforms, and so on. ■ Capture attributes or properties and relationships that are learned using NSX discovery agent, inventory collection, public cloud agent, Guest Introspection, VM Tools, and so on.
Security	■ Create grouping membership criteria. ■ Specify the firewall source and destination.
Troubleshooting (Traceability)	■ Trace a firewall rule into the logs (Rule tags) ■ Trace and correlate objects back to an OpenStack network.

System Tags

System tags are tags that are system-defined, and you cannot add, edit, or delete them.

Table 17-1. System Tags in Public Cloud Manager Objects

Objects	System Tags
Logical Switch Node Logical Router Logical Router Uplink Port Static Route DHCP Profile Firewall Section Rule List	■ CrossCloud ■ CloudType ■ CloudScope ■ CloudRegion ■ CloudVpcId ■ PcmId ■ EntityType
NAT Rule	■ CrossCloud ■ CloudType ■ CloudScope ■ CloudRegion ■ CloudVpcId ■ PcmId ■ EntityType ■ DefaultSnatRule ■ DefaultLinkLocalSNatRule/Cloud-Public-IP ■ DefaultSiNatRule

Table 17-2. System Tags in Cloud Service Manager (CSM) Objects

Objects	System Tags
BFD Health Monitoring Profile Transport Zone Uplink Host Switch Profile Transport Node Edge Cluster	■ CrossCloud ■ CloudType ■ CloudScope ■ CloudRegion ■ CloudVpcId ■ PcmId ■ EntityType

Table 17-3. System Tags in NSX Cloud VMs

Tag Source	System Tags
Amazon	■ aws:account ■ aws:availabilityzone ■ aws:region ■ aws:vpc ■ aws:subnet ■ aws:transit_vpc
Microsoft Azure	■ azure:subscription_id ■ azure:region ■ azure:vm_rg ■ azure:vnet_name ■ azure:vnet_rg ■ azure:transit_vnet_name ■ azure:transit_vnet_rg

Table 17-4. System Tags in Other NSX Objects

Objects	System Tags
Group	■ autoPlumbing ■ abstractionPath ■ NLB-VIP_ID ■ NLB-Lb-ID ■ NLB-Pool_ID
Segment	■ subnet-cidr
IP Address Pool IP Address Block	■ abstractionPath

Discovered Tags

NSX can discover and synchronize tags from the following:

- Amazon Web Services
- Microsoft Azure
- Kubernetes container clusters with NSX Container Plugin (NCP)
- Kubernetes container clusters with Antrea network plug-in
- OpenShift container clusters with NCP

Discovered tags are displayed for workload VMs and container cluster objects in the NSX Manager inventory. You cannot edit discovered tags in the UI. For example:

- Tags that are added to VMs in the public cloud and are automatically discovered by NSX Cloud. When you make changes to the tags in the public cloud, the changes are reflected in NSX Manager. By default, this feature is enabled.

- Tags that are added to container cluster objects in the Kubernetes container clusters with NCP CNI or Antrea CNI, and OpenShift container clusters with NCP CNI.

Tag Prefix	Meaning
dis:aws	Tags discovered from Amazon Web Services (AWS).
dis:azure	Tags discovered from Microsoft Azure.
dis:k8s	Tags discovered from Kubernetes container clusters with NCP CNI, Antrea CNI, and OpenShift container clusters with NCP CNI.

You can enable or disable the discovery of AWS tags at the time of adding the AWS account. Similarly, you can enable or disable Microsoft Azure tags at the time of adding the Microsoft Azure subscription.

Add Tags to an Object

You can select existing tags that are available in the NSX inventory or create new tags to add to an object.

The following procedure explains the steps for adding tags to a single object. For this procedure, the virtual machine object is considered. The steps for adding tags to other objects remain the same. You can navigate to the specific object page, and follow similar steps to add tags to that object.

For information about the maximum number of tags supported in NSX objects, see the VMware Configuration Maximums tool at https://configmax.vmware.com/home.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Edit an object for tags. Objects can be virtual machines, segments, or segment ports. You can use inventory for each object to tag the object or go to **Inventory > Tags** for creating and assigning tags.

   To tag the objects directly, for example virtual machines, click **Inventory > Virtual Machines**.

   Next to the virtual machine that you want to edit, click  , and click **Edit**.

3 In the **Tag** drop-down menu, enter a tag name. When you are done, click **Add Item(s)**.

   The maximum length of the tag name is 256 characters.

   If tags exist in the inventory, the **Tag** drop-down menu displays a list of all the available tags and their scope. The list of available tags includes user-defined tags, system-defined tags, and discovered tags. You can select an existing tag from the drop-down menu and add it to the virtual machine.

4 (Optional) Enter a tag scope.

   For example, let us say, you want to tag virtual machines based on their operating system (Windows, Mac, Linux). Create three tags, such as Windows, Linux, and Mac, and set the scope of each tag to OS.

The maximum length of the scope is 128 characters.

If you selected an existing tag from the inventory, the scope of the selected tag is applied automatically. Otherwise, you can enter a scope for the new tag that you are creating.

5 Click the **+** icon.

The tag is added to the virtual machine.

6 (Optional) Repeat steps 3–5 to add more tags to the virtual machine.

7 Click **Save**.

Add a Tag to Multiple Objects

You can add a tag to multiple objects simultaneously. However, this feature is available only for the virtual machine object.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Click **Inventory > Tags**.

3 Click **Add Tag**.

4 Enter a tag name.

The maximum length of the tag name is 256 characters.

5 (Optional) Enter a tag scope.

For example, let us say, you want to tag virtual machines based on their operating system (Windows, Mac, Linux). Create three tags, such as Windows, Linux, and Mac, and set the scope of each tag to OS.

The maximum length of the scope is 128 characters.

6 In **Assigned To**, click **Set Virtual Machines**.

7 (Required) Select one or more virtual machines to which you want to assign the tag, and click **Apply**.

You must assign a tag to at least one virtual machine before you can save the tag.

Note You can do a bulk assignment of a tag on a maximum of 1000 virtual machines at one go.

8 Click **Save**.

Results

- If the tag is assigned to many virtual machines, the assignment might take some time. When the assignment is in progress, the **Last Assignment Status** shows `Running`. After the tag is assigned successfully to all the selected virtual machines, the **Last Assignment Status** column changes to `Successful`.

- If a partial assignment occurs, NSX does not roll back the tag assignment from the VMs on which the tag is applied. For example, assume that you selected 100 VMs for a bulk tag assignment, and the assignment fails for 10 VMs. The tag that is assigned on the remaining 90 VMs is not rolled back.

  In such partial assignment situations, run the following API to retrieve the status of the tag operation:

  ```
  GET /api/v1/infra/tags/tag-operations/<tag-operation-id>/status
  ```

  You can also retrieve the realized status of the tag operation with the following API:

  ```
  GET /api/v1/infra/realized-state/realized-entities?intent_path=/infra/
  tags/tag-operations/<operation-id>
  ```

For more details about these APIs, see the *NSX API Guide*.

What to do next

If you have a long list of tags in the inventory, you can filter or search tags to find the tags of your interest quickly. You can filter on source, scope, and tag (name of the tag). You can also sort tags in the UI. However, due to the case-sensitive nature of tags, tags are sorted only in a lexical order.

The following limitations apply to searching or filtering tags:

- You cannot filter tags on the source and scope attributes simultaneously because both work on the scope attribute of the tag.
- The API does not support filtering tags with special characters, such as *, &, /, \, and so on. However, you can use special characters to filter tags in the UI.

Unassign Tags from an Object

You can remove tags that you had assigned previously to an object.

The following procedure explains the steps for unassigning tags from a single NSX object. For this procedure, the virtual machine object is considered. The steps for unassigning tags from other objects remain the same. You can navigate to the specific object page, and follow similar steps to unassign tags from that object.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Edit an object.

   For example, click **Inventory > Virtual Machines**. Next to the virtual machine that you want to edit, click the vertical ellipses, and click **Edit**.
3. Click the **X** icon for each tag that you want to unassign from the virtual machine.
4. Click **Save**.

Unassign a Tag from Multiple Objects

You can unassign a tag from multiple objects simultaneously. However, this feature is available only for virtual machines.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Click **Inventory > Tags**.

3 Next to that tag that you want to edit, click the vertical ellipses, and click **Edit**.

4 In the **Assigned To** column, click the number of virtual machines that are assigned this tag.

5 Click the **X** icon for each virtual machine that you want to unassign this tag.

Note

- You can do a bulk unassignment of a tag on a maximum of 1000 virtual machines at one go.
- When a tag is unassigned from all objects, the tag is deleted automatically from the inventory after five days.

6 Click **Apply**, and then click **Save**.

Resource Sharing

NSX multi-tenancy supports the sharing of certain objects (resources) from the **Default** view (default space) to specific projects or to all projects in the org.

You must have any one of these system-wide user roles to share objects from the default space:

- Enterprise Admin
- Network Admin
- Security Admin

Overview of Default Share

When an enterprise administrator adds a project in your org, no user-created objects exist in that project. A new project has access to only the system-defined NSX objects that are shared by default through the default share. In other words, the default share is created automatically when NSX is deployed. The objects in a default share are available to all the projects in the org.

The default share includes system-defined objects, such as, Services, BFD Profiles, App IDs, and many more.

To view the full list of system-defined objects in a default share, do these steps:

1 Ensure that you have selected the **Default** view in the project switcher drop-down menu.

2 Navigate to **Inventory > Resource Sharing**.

3 Expand the **Default Share**.

4 Click the number in the **Members** column.

For example:



Observe that in addition to the Default Share, which is shared with all projects, the system also automatically creates a default share for each project that is added in the system. The naming convention of the project-specific default share is "default-*<project-name>*".

The members of the project-specific default share are:

- Tier-0 Gateways (if set during project creation)
- Edge Clusters (if set during project creation)
- Site (if edge cluster is set during project creation)
- Site Enforcement Point (if edge cluster is set during project creation)

If neither tier-0/VRF gateways and edge clusters are selected during project creation, the project-specific default share has no members. These objects are managed from the default space and cannot be edited from the project-specific default share.

Overview of Resource Share

As mentioned earlier in this document, the following system-wide user roles can share objects from the default space to specific projects or to all the projects in the org:

- Enterprise Admin
- Network Admin
- Security Admin

This object sharing is done by creating resource shares. Each resource share is identified by a unique name. Project users can consume the shared objects in their projects to configure groups, firewall rules, and so on, to meet their networking and security requirements.

Shared objects are available in a **read-only** mode to the projects with which they are shared. In other words, shared objects cannot be modified by users in those projects.

Currently, system does not support sharing of objects from one project to another project. Objects can be shared only from the default space to other projects.

To view the shared objects in a project view, you must ensure that the **Shared Objects** check box is selected on the object listing page. For example, to view the shared segments in your project, switch to your project view, navigate to the **Segments** page, and at the bottom of this page, click the **Shared Objects** check box.

The following objects are currently supported for adding in a resource share:

- Groups
- Services
- Context Profiles
- Segments
- DHCP Profiles
- DAD Profiles
- ND Profiles

When you share an object from the default space by creating a resource share, the child objects of that shared object are not shared with the target project.

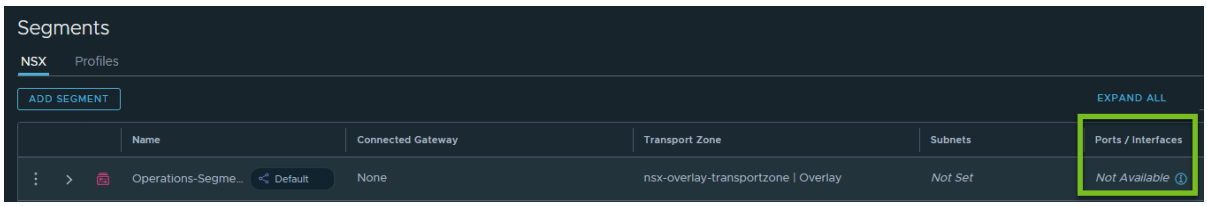
Use Case for Sharing Segments with Projects

When a segment is shared with a project, it does not mean that the VMs, ports, and gateway interfaces on this segment are exposed to the project. In fact, the project does not have visibility into the segment ports, interfaces, and workload VMs of the shared segment. Therefore, project users cannot configure distributed firewall policies on the workload VMs that are connected to the shared segment.

Sharing a segment with a project allows a project user to connect the segment to the service interface of a tier-1 gateway in that project.

Example: Consider the following scenario:

- Assume that in the **Default** space, there is an isolated segment named "Operations-Segment" and a tier-0 gateway named "T-0-Operations".
- On this tier-0 gateway, add a service interface and connect this interface to the "Operations-Segment".
- The enterprise administrator adds this segment to a resource share and shares it with project-1.
- Now, switch to project-1 in NSX Manager, navigate to the **Segments** page, and click the **Shared Objects** check box at the bottom of the page.
- Observe the properties of the shared "Operations-Segment". The **Ports/Interfaces** column shows value as `Not Available`. In other words, the service interface is not exposed to project-1.



Use Case for Sharing Groups, Services, Context Profiles with Projects

Typically, project users might want to consume NSX objects such as Groups, Services, and Context Profiles that exist in the default space of the system to create firewall rules under their projects. Resource sharing avoids the need for project users to recreate these objects. The shared objects become available to the projects in a **read only** mode, and they cannot be edited inside projects.

Share Resources from the Default View

You can share certain objects (resources) from the **Default** view (default space) to specific projects or to all the projects in the org. Shared objects are available in a **read-only** mode in the projects with which they are shared.

Prerequisites

You must have any one of these system-wide user roles to share objects from the default space:

- Enterprise Admin
- Network Admin
- Security Admin

Ensure that you have understood the concept of sharing resources. See Resource Sharing.

At least one project is added in your org. See Add a Project.

Procedure

1. From your browser, log in to an NSX Manager at `https://nsx-manager-ip-address`.
2. Navigate to **Inventory > Resource Sharing**.

   The **Resource Sharing** page is displayed. When you navigate to this page for the first time in your NSX deployment, only the system-defined **Default Share** is listed on this page. The members (objects) of the **Default Share** are available to all the projects in your org.
3. To create a resource share, click **Add Resource Share**.
4. Configure the properties of the resource share, as follows:
   a. Enter a name for the resource share.
   b. Click **Set**.

c In the **Set Members** window, click the object to be added in your resource share, and then select the object names.

Note Before adding segments in your resource share, make sure that you understand the typical use cases for sharing segments. For more information, see Resource Sharing.

d Click **Apply**.

e In the **Shared With** column, click **Set**, and then choose any one of these options:

Option	Description
Share with all Projects	Select this option to share the objects in this resource share with all the projects in your org.
Share with selected Projects	Select this option to share the objects in this resource share with specific projects.

f Click **Apply**.

5 To save the resource share, click **Save**.

Example: Resource Share

Let us assume that the enterprise administrator has shared a group from the default space with project-1 in your org.

To check whether this group is available for consumption in project-1, do these steps:

1 Log in to NSX Manager with the credentials of any user that has access to project-1.

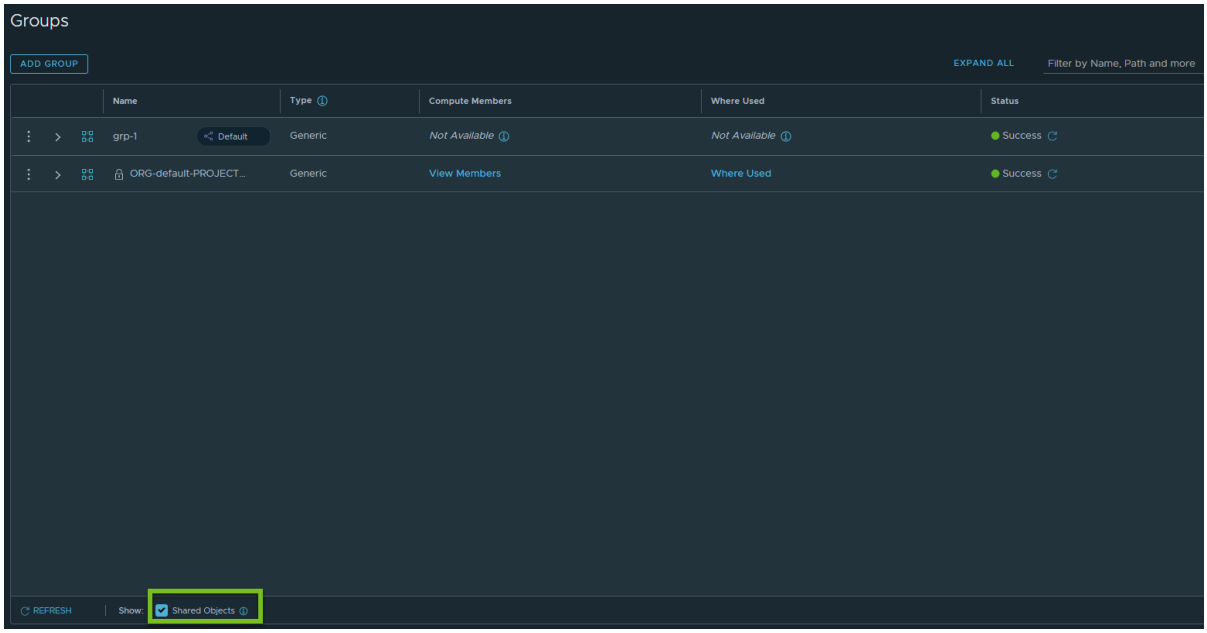
2 Ensure that **project-1** is selected from the project switcher drop-down menu.

3 Navigate to **Inventory > Groups**.

4 Click the **Shared Objects** check box at the bottom of the **Groups** page.

By default, this check box is not selected.

For example:



5 Observe that the shared group is listed on this page.

The following pill-shaped icon is displayed next to the object name to help you easily identify objects that are owned by the default space and shared with the project.



Multisite and NSX Federation 18

There are two options for managing NSX across multiple locations.

Table 18-1. Comparison of Multisite and NSX Federation

	Multisite	NSX Federation
Availability	NSX 2.3	NSX 3.0
Environments	Two locations in metropolitan regions (<10 ms across locations) with stretched VLAN	Other use cases
Number of NSX Manager clusters	1	1 per location
Network Services	All features supported: ■ Switching (Overlay and VLAN) ■ IPAM (DHCP and DNS) ■ Routing (VRF, EVPN, NAT and route redistribution) ■ Layer4+ services (Load Balancing, VPN)	All features supported from GM: ■ Switching (Overlay and VLAN) ■ IPAM (DHCP Relay and static binding, and DNS) ■ L2-Bridge ■ Routing (NAT and route redistribution) ■ Routing protocols (BGP, Static) Exceptions: ■ T0-VRF ■ DHCP dynamic binding ■ Routing protocols (OSPF) ■ Routing VPN and EVPN ■ Load Balancing
Security Services	All features supported: ■ Distributed Firewall ■ Gateway Firewall ■ FQDN Filtering ■ L7 App ID context support ■ Identity Firewall ■ Distributed IDS ■ Malware Prevention ■ Network Introspection ■ Endpoint Protection ■ Time-Based Firewall	All features supported from GM: ■ Distributed Firewall ■ Gateway Firewall ■ FQDN Filtering ■ L7 App ID context support ■ Time-Based Firewall Exceptions: ■ Identity Firewall ■ Distributed IDS ■ Malware Prevention ■ Network Introspection and Endpoint Protection

Table 18-1. Comparison of Multisite and NSX Federation (continued)

	Multisite	NSX Federation
High-Availability for Management Plane	NSX Manager VMs recovery ■ With SRM from NSX. See Working with VMware Site Recovery Manager and Multisite Environments. ■ With vSphere HA.	NSX GM and LM VMs recovery with SRM.
High-Availability for Compute VMs	Compute VMs recovery ■ With SRM from NSX. See Working with VMware Site Recovery Manager and Multisite Environments . ■ With vSphere HA.	Compute VMs recovery ■ With SRM from NSX with one limitation: Distributed Firewall on Workload VMs cannot be based on NSX tags. ■ With SRM from NSX without any limitations.

Also see the NSX-T Data Center Multi-location Design Guide and VMware Site Recovery Manager.

This chapter includes the following topics:

- NSX Multisite
- NSX Federation

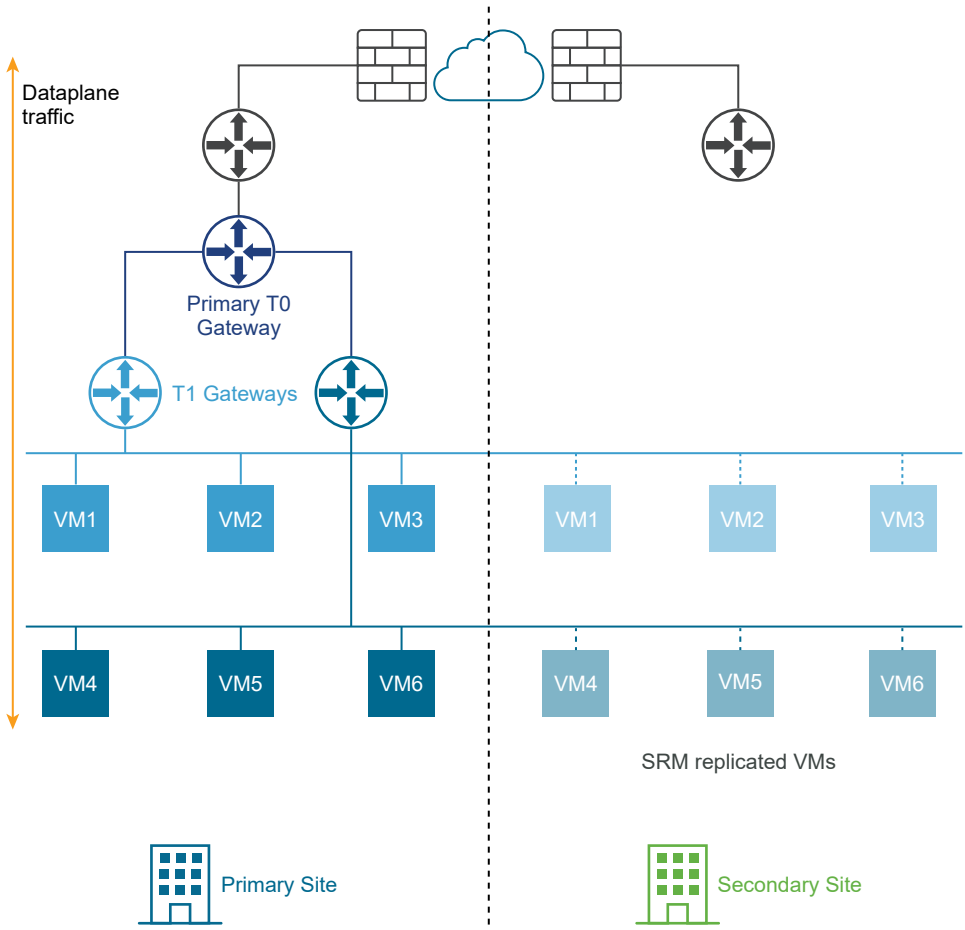
NSX Multisite

NSX supports multisite deployments where you can manage all the sites from one NSX Manager cluster.

Two types of multisite deployments are supported:

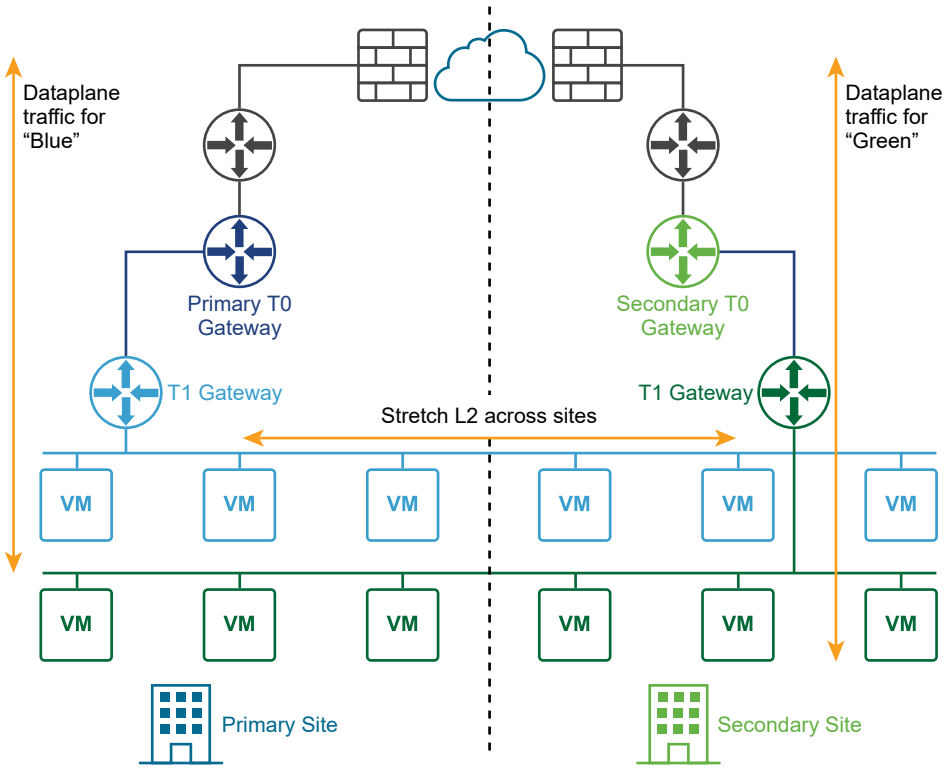
- Disaster recovery
- Active-active

The following diagram illustrates a disaster recovery deployment.



In a disaster recovery deployment, NSX at the primary site handles networking for the enterprise. The secondary site stands by to take over if a catastrophic failure occurs at the primary site.

The following diagram illustrates an active-active deployment.



You can deploy two sites for automatic or manual/scripted recovery of the management plane and the data plane.

Automatic Recovery of the Management Plane

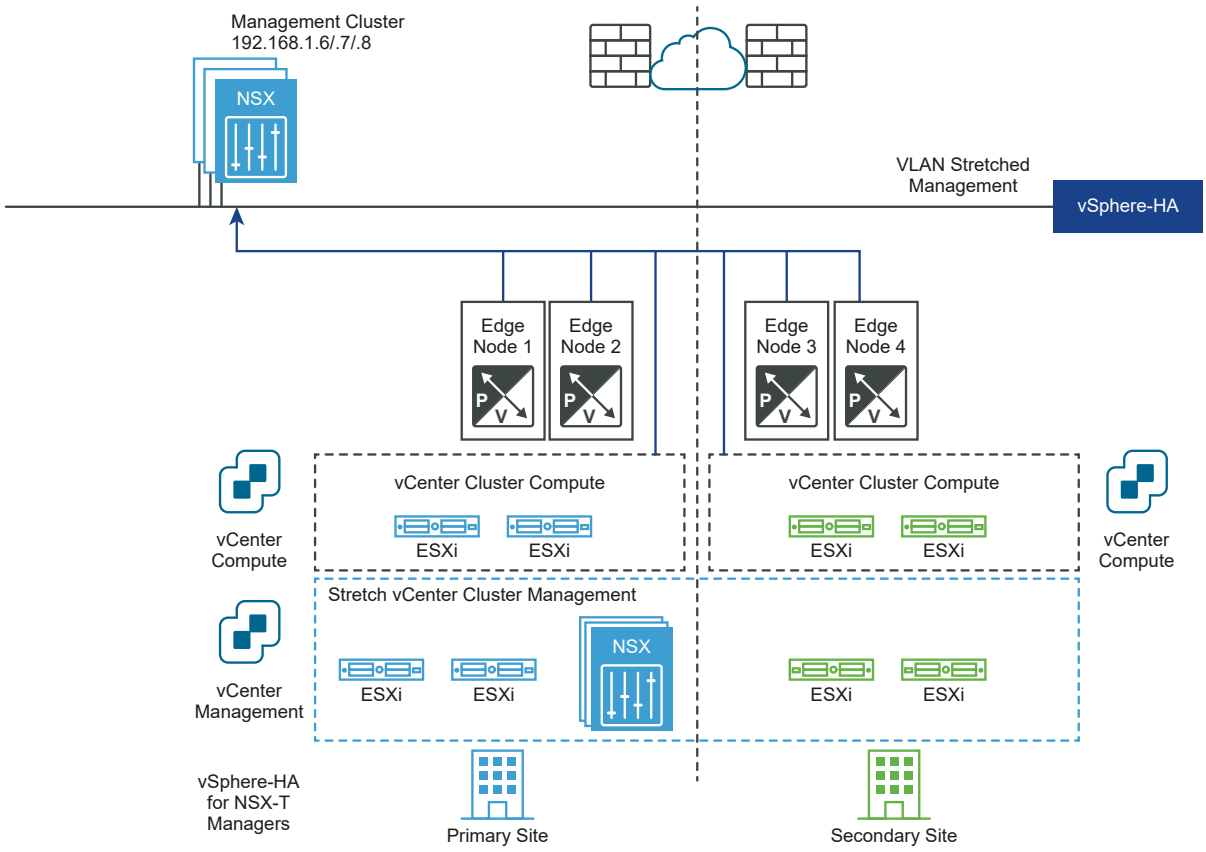
Requirements:

- A stretched vCenter cluster with HA across sites configured.
- A stretched management VLAN.

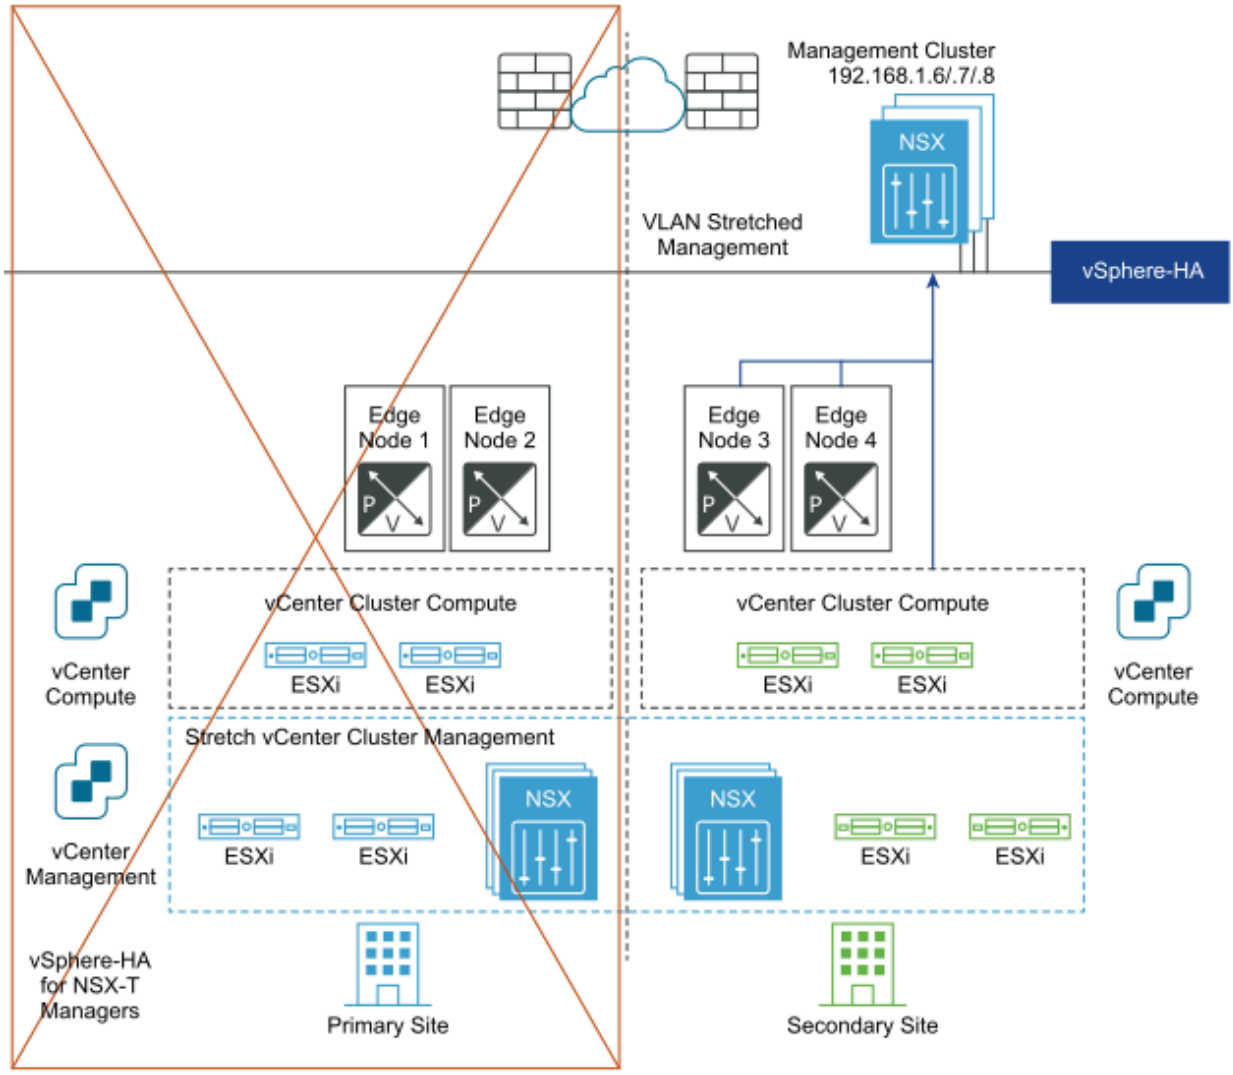
The NSX Manager cluster is deployed on the management VLAN and is physically in the primary site. If there is a primary site failure, vSphere HA will restart the NSX Managers in the secondary site. All the transport nodes will reconnect to the restarted NSX Managers automatically. This process takes about 10 minutes. During this time, the management plane is not available but the data plane is not impacted.

The following diagrams illustrates automatic recovery of the management plane.

Before the disaster:



After disaster recovery:



Automatic Recovery of the Data Plane

You can configure failure domains for Edge nodes to achieve automatic recovery of the data plane. You can group Edge nodes within an Edge cluster in different failure domains. NSX Manager will automatically place any new active tier-1 gateway in the preferred failure domain, and the standby tier-1 gateway in the other domain. Note that a T1 deployed prior to the failure domain creations will keep their original Edge Node placement and might not be running where you want. If you want to fix their placement, edit the T1 and manually select the Edge Nodes for T1-Active and T1-Standby.

Requirements:

- The maximum latency between Edge nodes is 10 ms.
- If asymmetric north-south routing is not achievable, for example a physical firewall is used northbound to the NSX Edge node, then the HA mode for the tier-0 gateway must be active-standby, and the failover mode must be preemptive.

- If asymmetric north-south routing is possible, for example the two locations are two buildings without any physical firewall between them, then the HA mode for the tier-0 gateway can be active-active.

The Edge nodes can be VMs or bare metal. The failover mode of the tier-1 gateway can be preemptive or non-preemptive, but preemptive is recommended to guarantee that the tier-0 and tier-1 gateways are in the same location.

Configuration steps:

- Using the API, create failure domains for the two sites, for example, `FD1A-Preferred_Site1` and `FD2A-Preferred_Site1`. Set the parameter `preferred_active_edge_services` to `true` for the primary site and set it to `false` for the secondary site.

```
POST /api/v1/failure-domains
{
"display_name": "FD1A-Preferred_Site1",
"preferred_active_edge_services": "true"
}

POST /api/v1/failure-domains
{
"display_name": "FD2A-Preferred_Site1",
"preferred_active_edge_services": "false"
}
```

- Using the API, configure an Edge cluster that is stretched across the two sites. For example, the cluster has Edge nodes `EdgeNode1A` and `EdgeNode1B` in the primary site, and Edge nodes `EdgeNode2A` and `EdgeNode2B` in the secondary site. The active tier-0 and tier-1 gateways will run on `EdgeNode1A` and `EdgeNode1B`. The standby tier-0 and tier-1 gateways will run on `EdgeNode2A` and `EdgeNode2B`.
- Using the API, associate each Edge node with the failure domain for the site. First call the `GET /api/v1/transport-nodes/<transport-node-id>` API to get the data about the Edge node. Use the result of the GET API as the input for the `PUT /api/v1/transport-nodes/<transport-node-id>` API, with the additional property, `failure_domain_id`, set appropriately. For example,

```
GET /api/v1/transport-nodes/<transport-node-id>
Response:
{
    "resource_type": "TransportNode",
    "description": "Updated NSX configured Test Transport Node",
    "id": "77816de2-39c3-436c-b891-54d31f580961",
    ...
}

PUT /api/v1/transport-nodes/<transport-node-id>
{
    "resource_type": "TransportNode",
    "description": "Updated NSX configured Test Transport Node",
```

```
    "id": "77816de2-39c3-436c-b891-54d31f580961",
    ...
    "failure_domain_id": "<UUID>",
}
```

- Using the API, configure the Edge cluster to allocate nodes based on failure domain. First call the `GET /api/v1/edge-clusters/<edge-cluster-id>` API to get the data about the Edge cluster. Use the result of the GET API as the input for the `PUT /api/v1/edge-clusters/<edge-cluster-id>` API, with the additional property, `allocation_rules`, set appropriately. For example,

```
GET /api/v1/edge-clusters/<edge-cluster-id>
Response:
{
    "_revision": 0,
    "id": "bf8d4daf-93f6-4c23-af38-63f6d372e14e",
    "resource_type": "EdgeCluster",
    ...
}

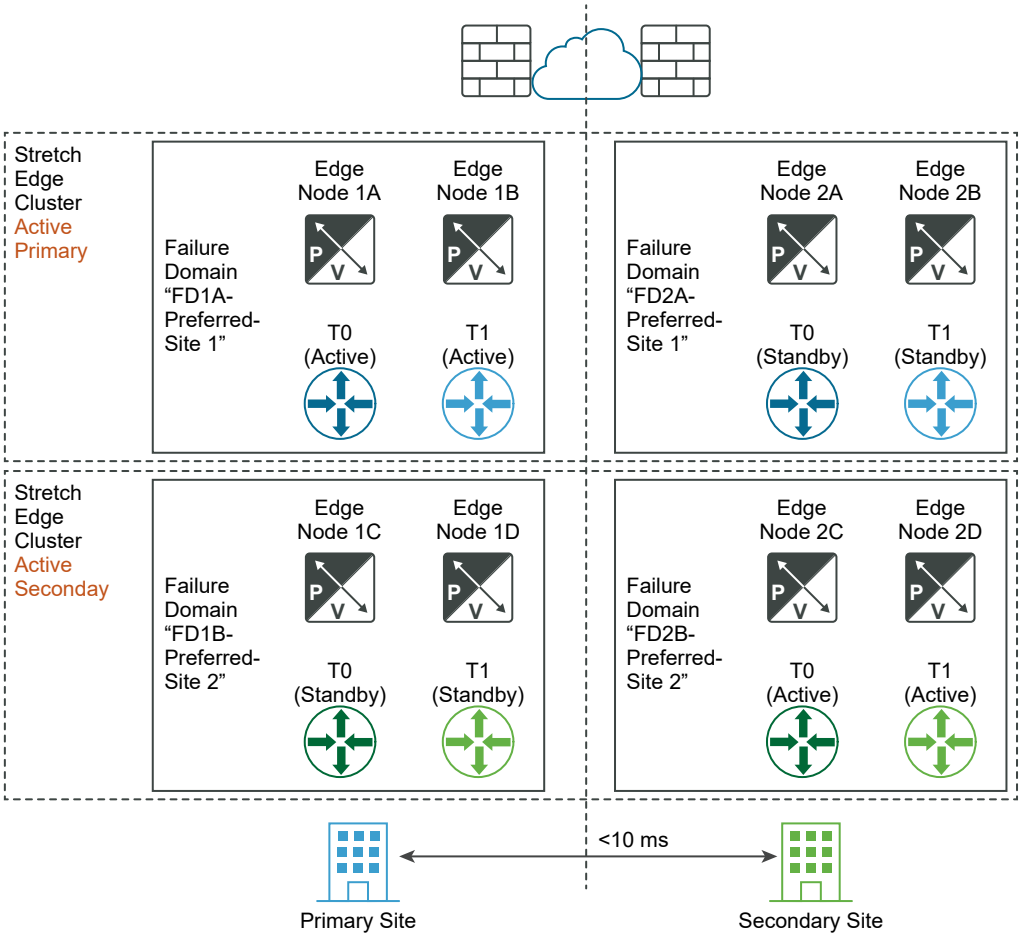
PUT /api/v1/edge-clusters/<edge-cluster-id>
{
    "_revision": 0,
    "id": "bf8d4daf-93f6-4c23-af38-63f6d372e14e",
    "resource_type": "EdgeCluster",
    ...
    "allocation_rules": [
        {
            "action": {
                      "enabled": true,
                      "action_type": "AllocationBasedOnFailureDomain"
                      }
        }
    ],
}
```

- Create tier-0 and tier-1 gateways using the API or NSX Manager UI.

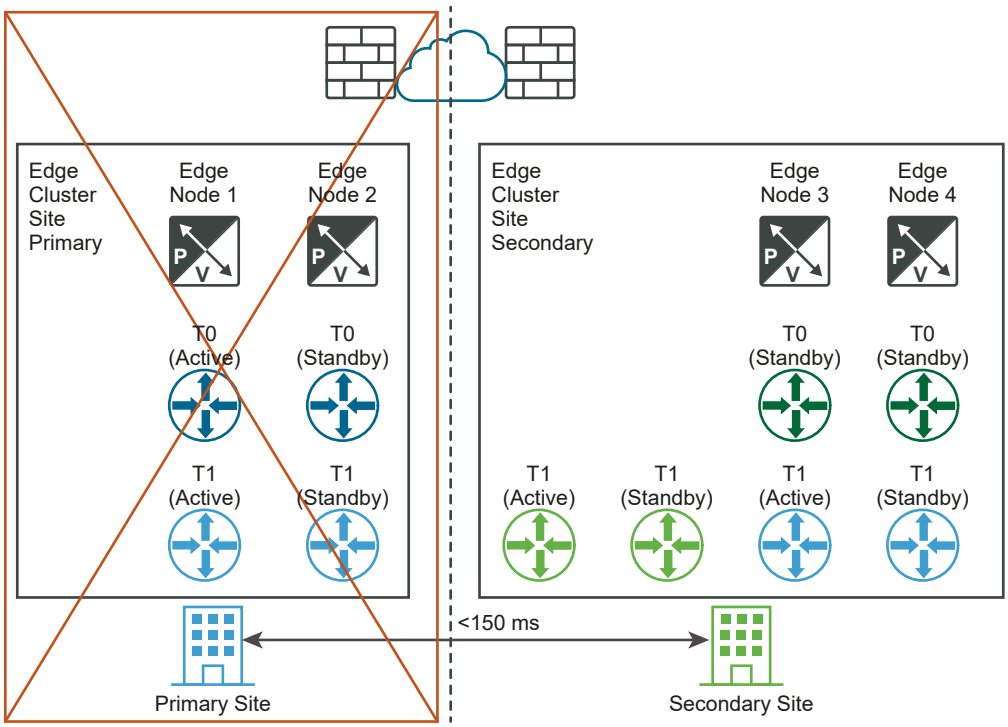
In case of a full primary site failure, the tier-0 standby and tier-1 standby in the secondary site automatically take over and become the new active gateways.

The following diagrams illustrates automatic recovery of the data plane.

Before the disaster:



After disaster recovery:



In case of a failure of one of the Edge nodes in the primary site and not full site failure, it is important to note that the same principle applies. For example, in the diagram, "Before the disaster", assume that Edge node 1B hosts tier-1-blue active and that Edge node 2B hosts the tier-1-blue standby. If Edge node 1B fails, the standby tier-1-blue on Edge node 2B takes over and become the new tier-1-blue active gateway.

Manual/Scripted Recovery of the Management Plane

Requirements:

- DNS for NSX Manager with a short TTL (for example, 5 minutes).
- Continuous NSX Manager backup.

Neither vSphere HA, nor a stretched management VLAN, is required. NSX Managers must be associated with a DNS name with a short TTL. All transport nodes (Edge nodes and hypervisors) must connect to the NSX Manager using their DNS name. To save time, you can optionally pre-install an NSX Manager cluster in the secondary site.

The recovery steps are:

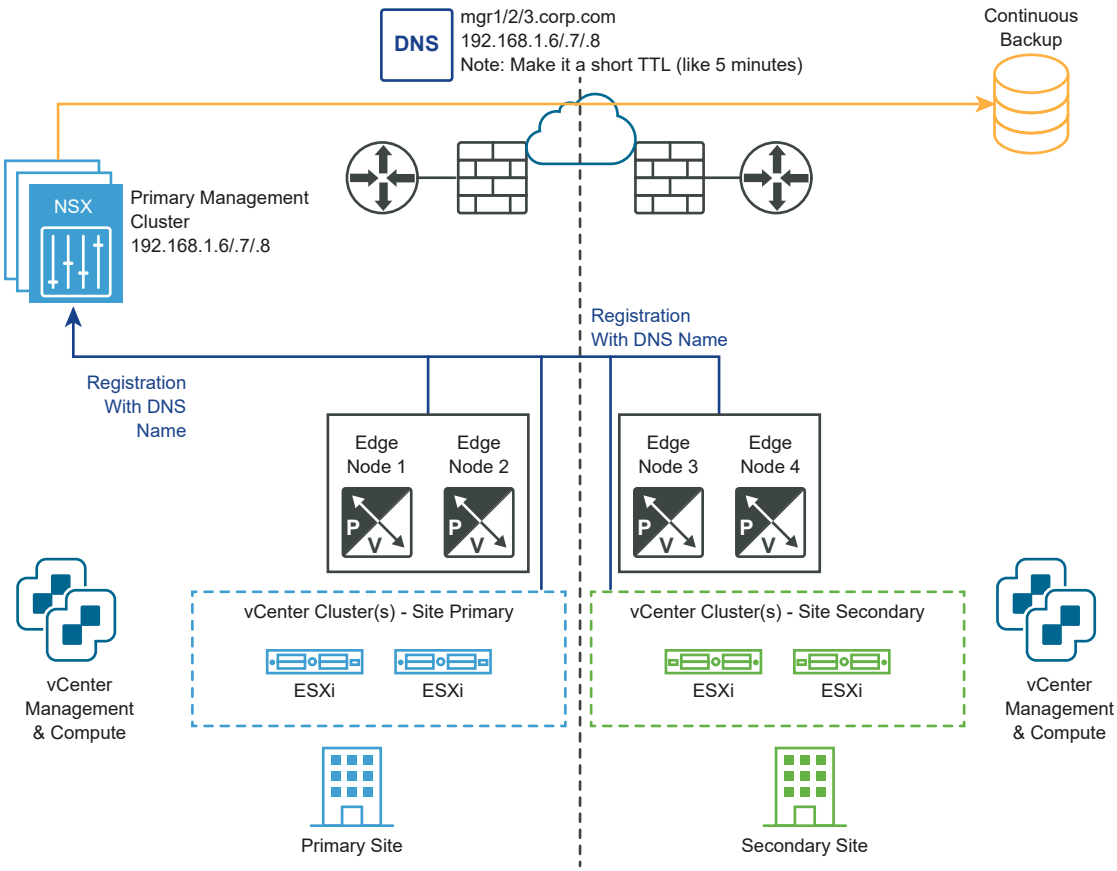
1 Change the DNS record so that the NSX Manager cluster has different IP addresses.

2 Restore the NSX Manager cluster from a backup.

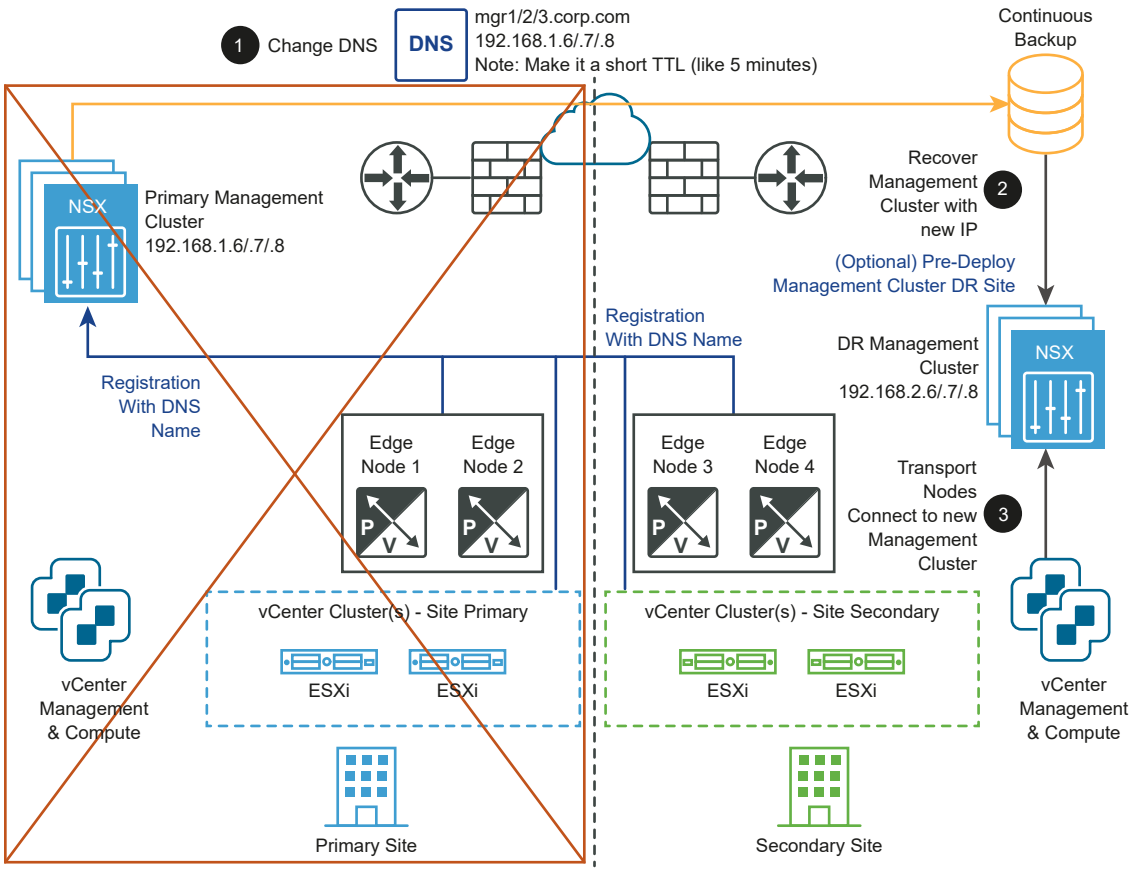
3 Connect the transport nodes to the new NSX Manager cluster.

The following diagrams illustrate manual/scripted recovery of the management plane.

Before the disaster:



After the disaster:



Manual/Scripted Recovery of the Data Plane

Requirement: The maximum latency between Edge nodes is 150 ms.

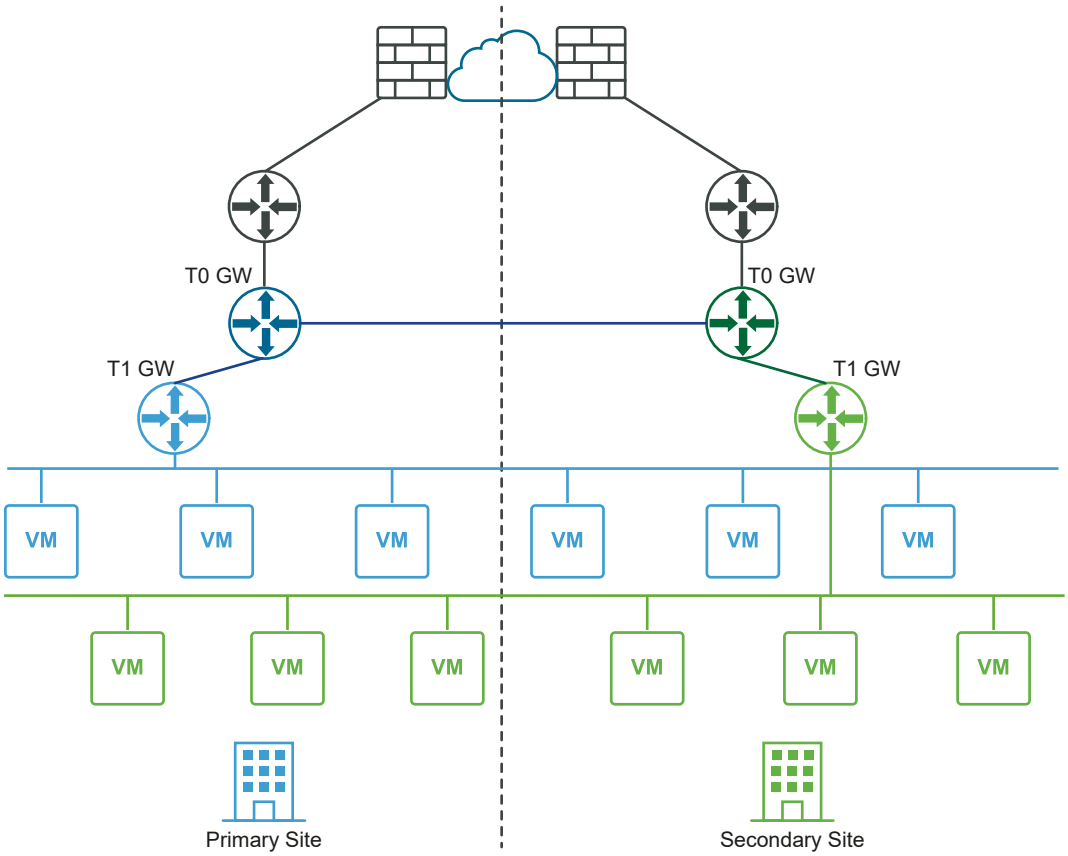
The Edge nodes can be VMs or bare metal. The tier-0 gateways in each location can be active-standby or active-active. Edge node VMs can be installed in different vCenter Servers. No vSphere HA is required.

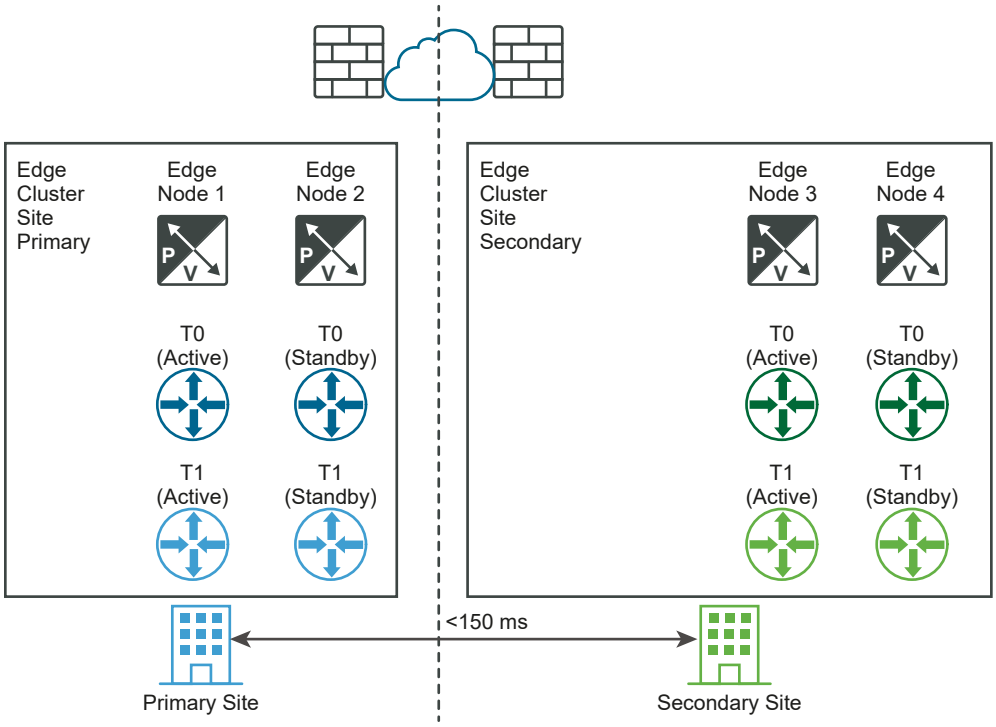
The recovery steps are:

- For all tier-1 in primary site (blue), update their Edge Cluster configuration to be the Edge Cluster Site secondary.
- For all tier-1 in primary site (blue), reconnect them to T0 secondary (green).

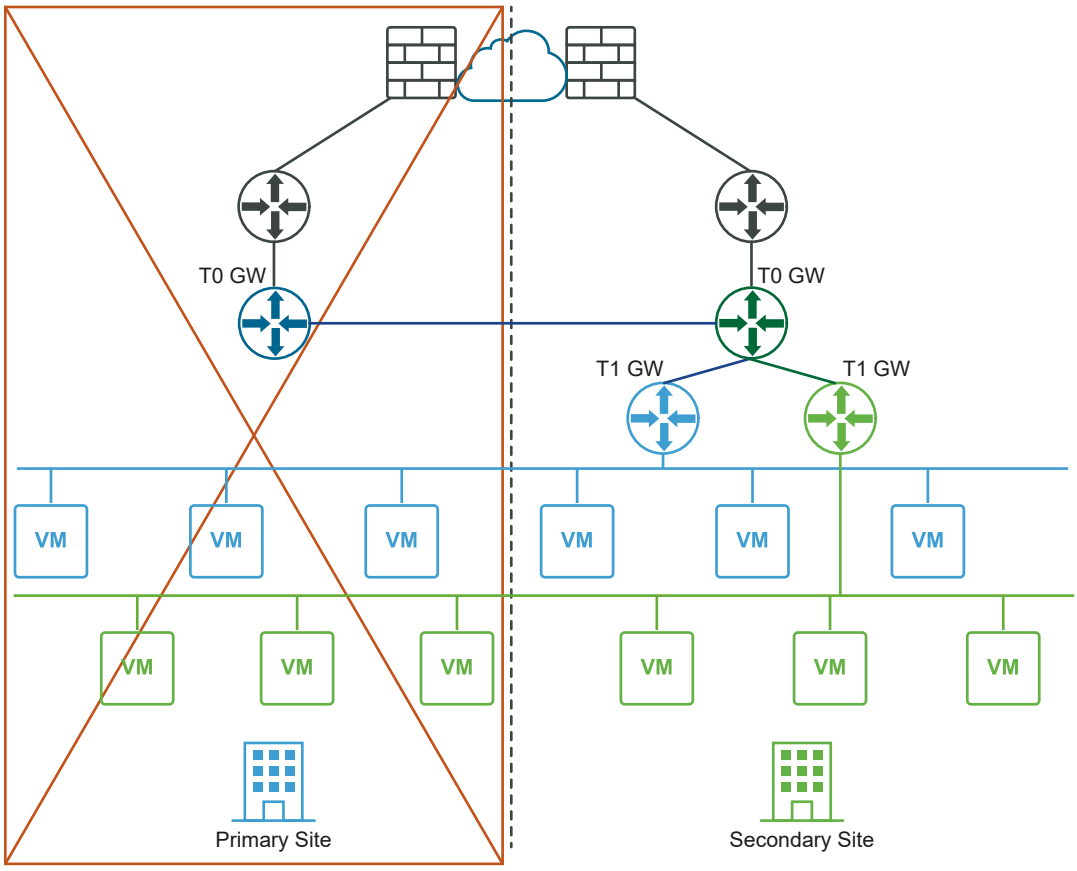
The following diagrams illustrate manual/scripted recovery of the data plane with both the logical and physical network views.

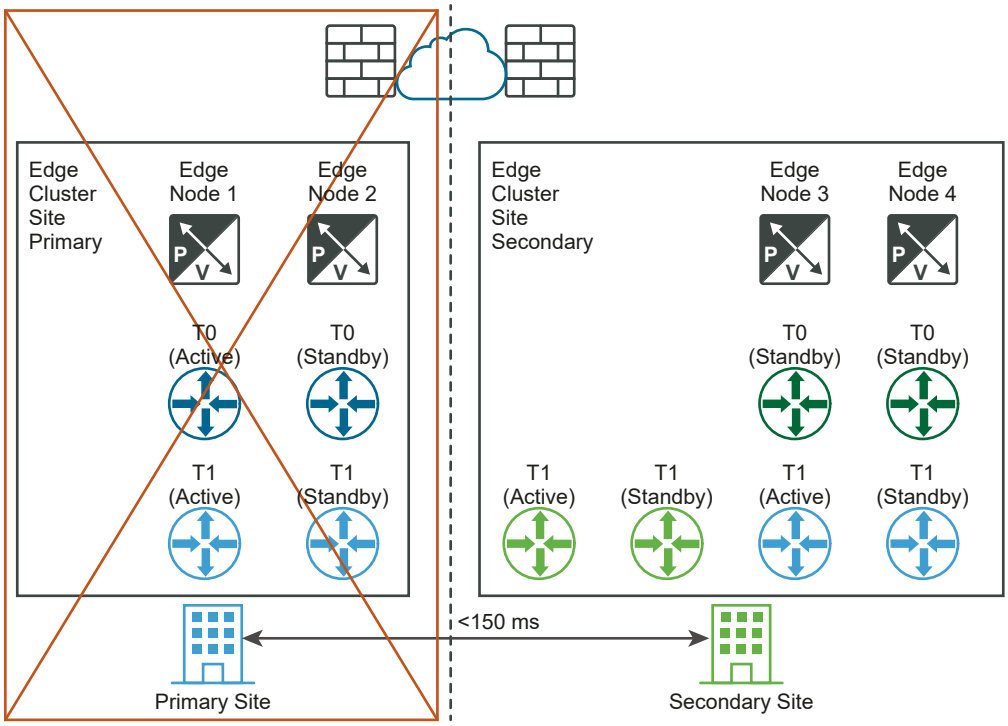
Before the disaster (logical and physical views):





After the disaster (logical and physical views):


T0 GW
T0 GW
T1 GW
T1 GW
VM
VM
VM
VM
VM
VM
VM
VM
VM
VM
VM
VM
Primary Site
Secondary Site


Edge Cluster Site Primary
Edge Node 1
Edge Node 2
P
V
T0 (Active)
T0 (Standby)
T1 (Active)
T1 (Standby)
Edge Cluster Site Secondary
Edge Node 3
Edge Node 4
T0 (Standby)
T0 (Standby)
T1 (Active)
T1 (Standby)
T1 (Active)
T1 (Standby)
<150 ms
Primary Site
Secondary Site

Requirements for Multisite Deployments

Inter-site Communication

- The bandwidth must be at least 1 Gbps and the latency (RTT) must be less than 150 ms.
- MTU must be at least 1600. 9000 is recommended.

NSX Manager

- With Automatic Recovery of Management Plane
  - VLAN management stretched between sites.
  - vSphere HA across sites for NSX Manager VMs.
- With Manual/Scripted Recovery of the Management Plane
  - Continuous NSX Manager backup.
  - NSX Manager must be set up to use FQDN.

Data Plane

- The same internet provider must be used if public IP addresses are exposed through services such as NAT or load balancer.
- With Automatic Recovery of Management Plane
  - Maximum latency between locations is 10 ms.
  - The HA mode for the tier-0 gateway must be active-standby and the failover mode must be preemptive to guarantee no asymmetric routing.
  - The HA mode for the tier-0 gateway can be active-active if asymmetric routing is acceptable (such as different buildings in a metropolitan region).
- With Manual/Scripted Recovery of Management Plane
  - Maximum latency between locations is 150 ms.

Cloud Management System

- The cloud management system (CMS) must support an NSX plug-in. In this release, VMware Integrated OpenStack (VIO) and vRealize Automation (vRA) satisfy this requirement.

Limitations

- No local-egress capabilities. All north-south traffic must occur within one site.
- The compute disaster recovery software must support NSX, for example, VMware SRM 8.1.2 or later.

- When restoring the NSX Manager in a multi-site environment do the following on the secondary/primary site:
  - After the restore process pauses at the **AddNodeToCluster** step, you must first remove the existing VIP and set the new virtual IP from the **System > Appliances** UI page before you add additional manager nodes.
  - Add new nodes to a restored one-node cluster after the VIP is updated.

Working with VMware Site Recovery Manager and Multisite Environments

You can use VMware Site Recovery Manager™ with NSX Multisite version 3.0.2 or later for disaster recovery use cases.

For detailed instructions on using Site Recovery Manager, see the VMware Site Recovery Manager Documentation.

Site Recovery Manager supports the following workflows with NSX Multisite:

- NSX Management VMs support full and test recovery of Management VMs (supported with or without NSX management cluster VIP).
- NSX Compute VMs support full and test recovery of compute VMs. Recovered VMs in the disaster recovery site have their NSX tags and any firewall rules based on these NSX tags or other VM attributes, such as VM name. If a VM gets deleted from ESXi, NSX removes the VM from its inventory but saves its VM tag information internally for 30 minutes. If that VM, with same instanceUuid, gets recovered in that time frame it gets its original VM tags. After 30 minutes, the NSX admin must reconfigure its VM tags. For VMs on ESXi, instanceUUid and external Id values are the same.

  Note For SRM planned migration:

  - SRM with array-based replication removes the VMs from the protected ESXi site and recovers those VMs to the recovery ESXi site. If that process takes more than 30 minutes, the VMs lose their NSX tags. As a result, the NSX admin must reconfigure those NSX tags.
  - SRM with vSphere Replication does not remove the VMs from the protected ESXi site. It powers them off and recover those VMs to the recovery ESXi site. As a result, the VMs never lose their VM tags.

For NSX Federation support of Site Recovery Manager, see Working with Site Recovery Manager and Federation.

NSX Federation

With NSX Federation, you can manage multiple NSX environments with a single pane of glass view, create gateways and segments that span one or more locations, and configure and enforce firewall rules consistently across locations.

Once you have installed the Global Manager and have added locations, you can configure networking and security from Global Manager.

For information about the initial NSX Federation configuration, including installing Global Manager and adding locations, see *Getting Started with NSX Federation* in the *NSX Installation Guide*.

Overview of NSX Federation

Before you configure your NSX Federation environment, understand what features are supported, how NSX Federation shares information across locations, and how the user interface works.

See NSX Federation Key Concepts for definitions of terminology specific to NSX Federation.

NSX Federation Key Concepts

NSX Federation introduces some new terms and concepts, such as remote tunnel endpoint (RTEP), span, and region.

NSX Federation Systems: Global Manager and Local Manager

An NSX Federation environment includes two types of management systems:

- Global Manager: A system similar to NSX Manager that federates multiple Local Managers.
- Local Manager: An NSX Manager system in charge of network and security services for a location.

NSX Federation Span: Local and Stretched

When you create a networking object from Global Manager, it can span one or more locations.

- Local: The object spans only one location.
- Stretched: The object spans more than one location.

You do not directly configure the span of a segment. A segment has the same span as the gateway it is attached to.

NSX Federation Regions

Security objects have a region. The region can be one of the following:

- Location: Each location automatically creates a region. This region has the span of that location.
- Global: A region that has the span of all available locations.
- Custom Region: You can create regions that include a subset of the available locations.

NSX Federation Tunnel Endpoints

In an NSX Federation environment, there are two types of tunnel endpoints.

- Tunnel End Point (TEP): The IP address of a transport node (Edge node or Host) used for Geneve encapsulation within a location.

- Remote Tunnel End Points (RTEP): The IP address of a transport node (Edge node only) used for Geneve encapsulation across locations.

Features and Configurations Supported in NSX Federation

To make all configurations from the Global Manager, use policy mode. The manager mode is not available in NSX Federation.

Refer to Chapter 1 NSX Manager for more information about the two modes.

Configuration Maximums

An NSX Federation environment has the following configuration maximums:

- For most configurations, the Local Manager cluster has the same configuration maximums as an NSX Manager cluster. Go to VMware Configuration Maximums tool and select NSX.

  Select the NSX Federation category for NSX in the VMware Configuration Maximums tool for exceptions and other NSX Federation-specific values.

- For a given location, the following configurations contribute to the configuration maximum:
  - Objects that were created on the Local Manager.
  - Objects that were created on the Global Manager and include the location in its span.

  You can view the capacity and usage on each Local Manager. See View the Usage and Capacity of Categories of Objects.

Feature Support

Service insertion (Network Introspection) support only occurs when an NSX Federation environment has a Global Manager (GM) deployed under the following conditions:

- All service-insertion related configuration such as partner service registration, deployment and consumption, is done from a Local Manager (LM).
- Only objects configured on the LM are used with service insertion. This includes groups, segments, and any other constructs. Service insertion cannot be applied to workloads connected to a stretched/global segment defined from the GM, or any segment connected to a logical router created from the GM. Groups created from the Global Manager should not be used within service insertion redirection polices.

Important

- NSX Federation locations must run on environments where administrators have full control of the underlay fabric.
- NSX Federation does not support Local Manager or Global Manager hosted on VMware Cloud on AWS (VMC on AWS), Azure VMware Solution (AVS), Google Cloud VMware Engine (GCVE), Oracle Cloud VMware Solution (OCVS), or Alibaba Cloud VMware Service (ACVS).

Table 18-2. Features Supported in NSX Federation

Feature	Details	Related Links
Tier-0 Gateway	■ Active-active and active-standby ■ Active-active only	Add a Tier-0 Gateway from Global Manager
Tier-1 Gateway		Add a Tier-1 Gateway from Global Manager
Segments	Includes Layer 2 bridge configuration from Global Manager.	Add a Segment from Global Managerand Configure Bridging on Global Manager
Groups	Some limitations. See Security in NSX Federation .	Create Groups from Global Manager
Distributed Firewall	Draft of the security policies are available on Global Manager. This includes support for auto and manual drafts.	Create Drafts In Global Manager
Firewall Exclusion List	Available.	Manage a Firewall Exclusion List
Time Based Firewall Rules	Available.	Time-Based Firewall Policy
Gateway Firewall		Create Gateway Policies and Rules from Global Manager
Network Address Translation (NAT)	Tier-0 Gateway: ■ Active-active: You can configure stateless NAT only, that is, with action type Reflexive. ■ Active-standby: You can create stateful or stateless NAT rules. Tier-1 Gateway: ■ You can create stateful or stateless NAT rules. Stateless NAT rules are pushed to all locations in the gateway's span unless scoped to one or more locations specifically. Stateful NAT rules are also pushed to all locations in the gateway's span or to the specific location selected. However, stateful NAT rules are realized and enforced only on the primary location.	Configure NAT/DNAT/No SNAT/No DNAT/Reflexive NAT
DNS		Add a DNS Forwarder Service

Table 18-2. Features Supported in NSX Federation (continued)

Feature	Details	Related Links

DHCP and SLAAC	■ DHCP Relay is supported on segments and gateways.
	■ DHCPv4 server is supported on gateways with DHCP static bindings configured on segments.
	■ IPv6 addresses can be assigned using SLAAC with DNS Through RA (DAD detects duplicates within a location only).
		■ DHCP Relay: Add a DHCP Relay Profile
		■ DHCP Server (supported on gateway only):
		  ■ Add a DHCP Server Profile
		  ■ Attach a DHCP Profile to a Tier-0 or Tier-1 Gateway
		  ■ DHCP Configuration Settings: Reference
		■ IPv6 address assignment: Create SLAAC and DAD Profiles for IPv6 Address Assignment

Using objects created on Global Manager in a Local Manager confiiguration	Most configurations are supported. For example: ■ Connecting a Local Manager tier-1 gateway to a Global Manager tier-0 gateway. ■ Using a Global Manager group in a Local Manager distributed firewall rule. These configurations are not supported: ■ Connecting a Local Manager segment to a Global Manager tier-0 or tier-1 gateway. ■ Connecting a load balancer to a Global Manager tier-1 gateway.	
Network Monitoring	■ Expanded communication monitoring between Local Manager and Global Manager. ■ Traceflow across NSX instances in the same federation.	
Certificates	Starting with NSX 4.1, Local Manager self-signed certificates generate only when the Local Manager is in the NSX Federation environment. That same certificate gets deleted if Local Manager moves out of the NSX Federation environment.	Certificates for NSX and NSX Federation
LDAP	Authenticate Global Managerusers using a directory service such as Active Directory over LDAP or OpenLDAP.	Integration with LDAP
Backup and Restore	■ Backup with FQDN or IP is supported.	Backup and Restore in NSX Federation

Table 18-2. Features Supported in NSX Federation (continued)

Feature	Details	Related Links
Users	Starting with NSX 4.1, you can add new local users and remove audit and guest users.	Manage Local User Accounts
vMotion between locations	■ Tag replication across locations is supported.	

Understanding NSX Federation

In NSX Federation, you make configuration changes on the active Global Manager. The active Global Manager then synchronizes the changes with the relevant Local Managers and the standby Global Manager, if you have one. Local Managers also sync some information with each other and to the Global Manager.

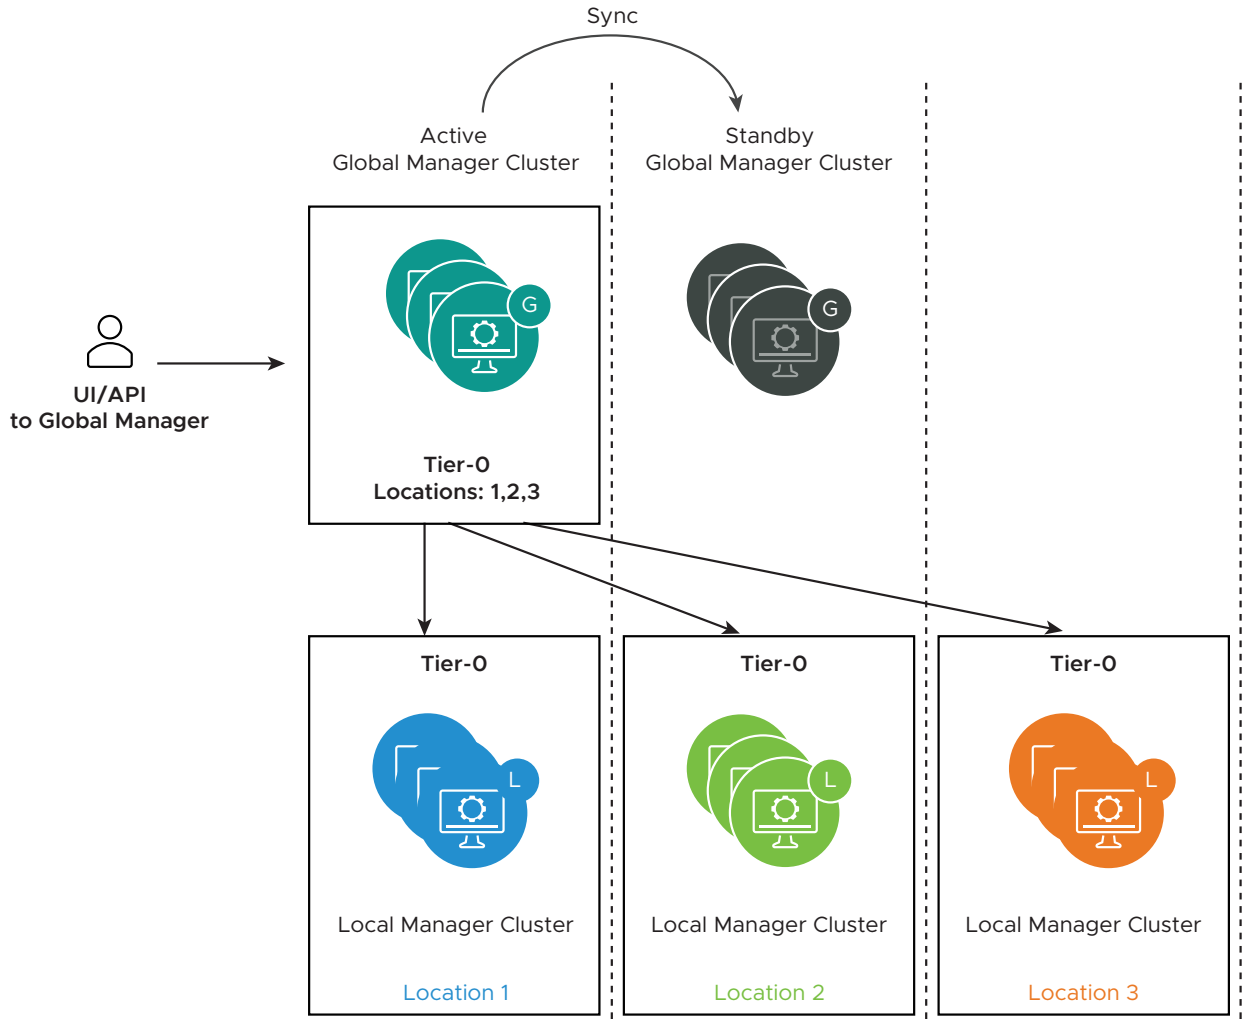
Making Changes on Global Manager

The Global Manager provides a user interface similar to the NSX Manager interface.

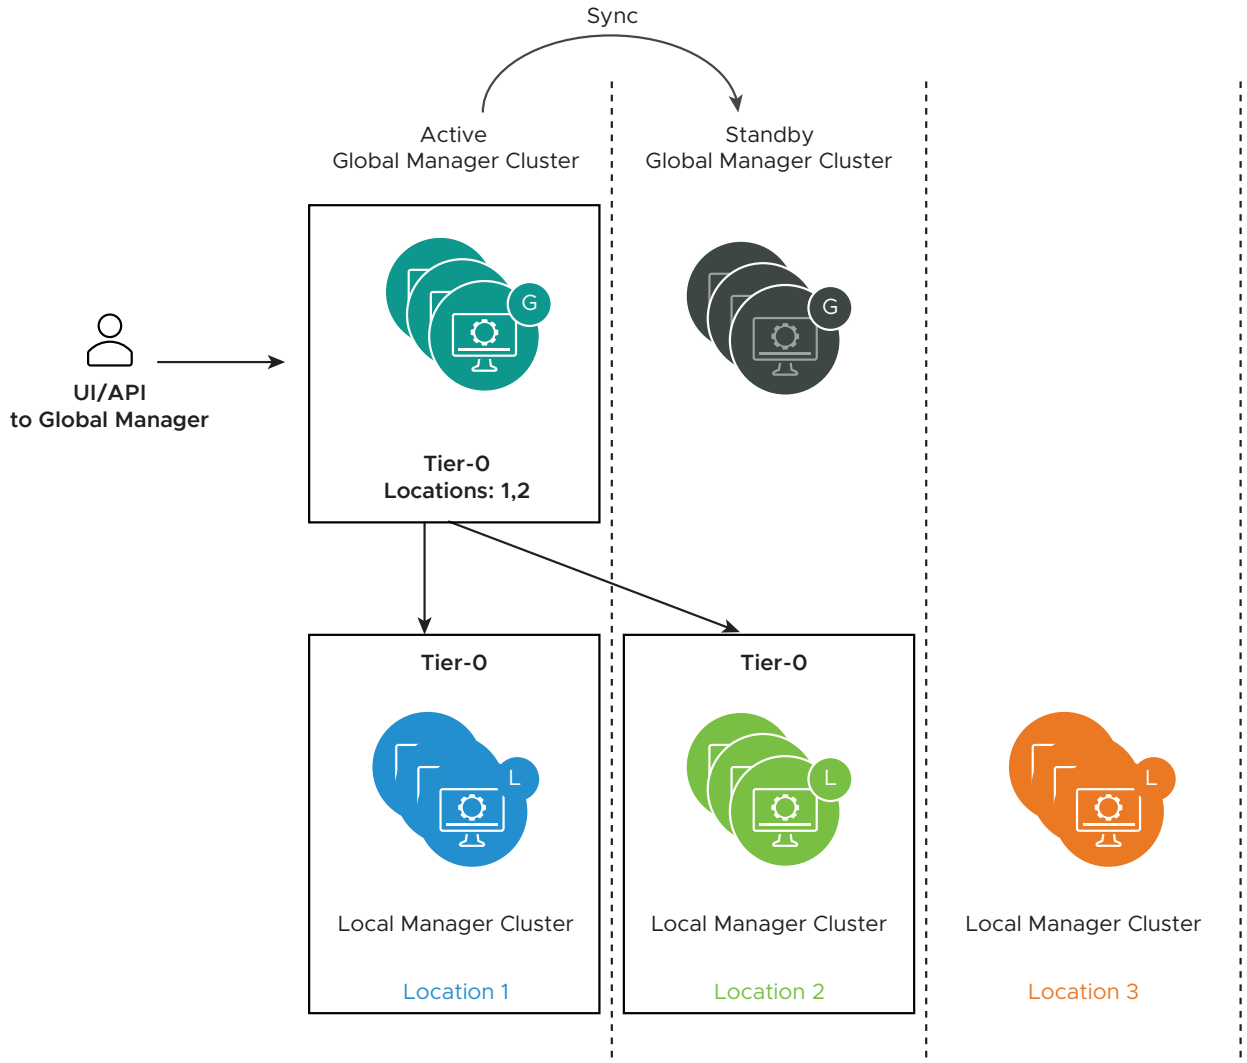
If you choose, you can configure all objects on the Global Manager, independent of span.

The Global Manager syncs a configuration with a Local Manager only if the configuration is relevant to that location. For example, if you create a tier-0 gateway and add it to Location 1, Location 2, and Location 3, the configuration gets synchronized with all three Local Managers. Local Managers can only synchronize once with Global Manager during a configuration import.

If you have a standby Global Manager, the configurations synchronize between the active Global Manager and the standby Global Manager.

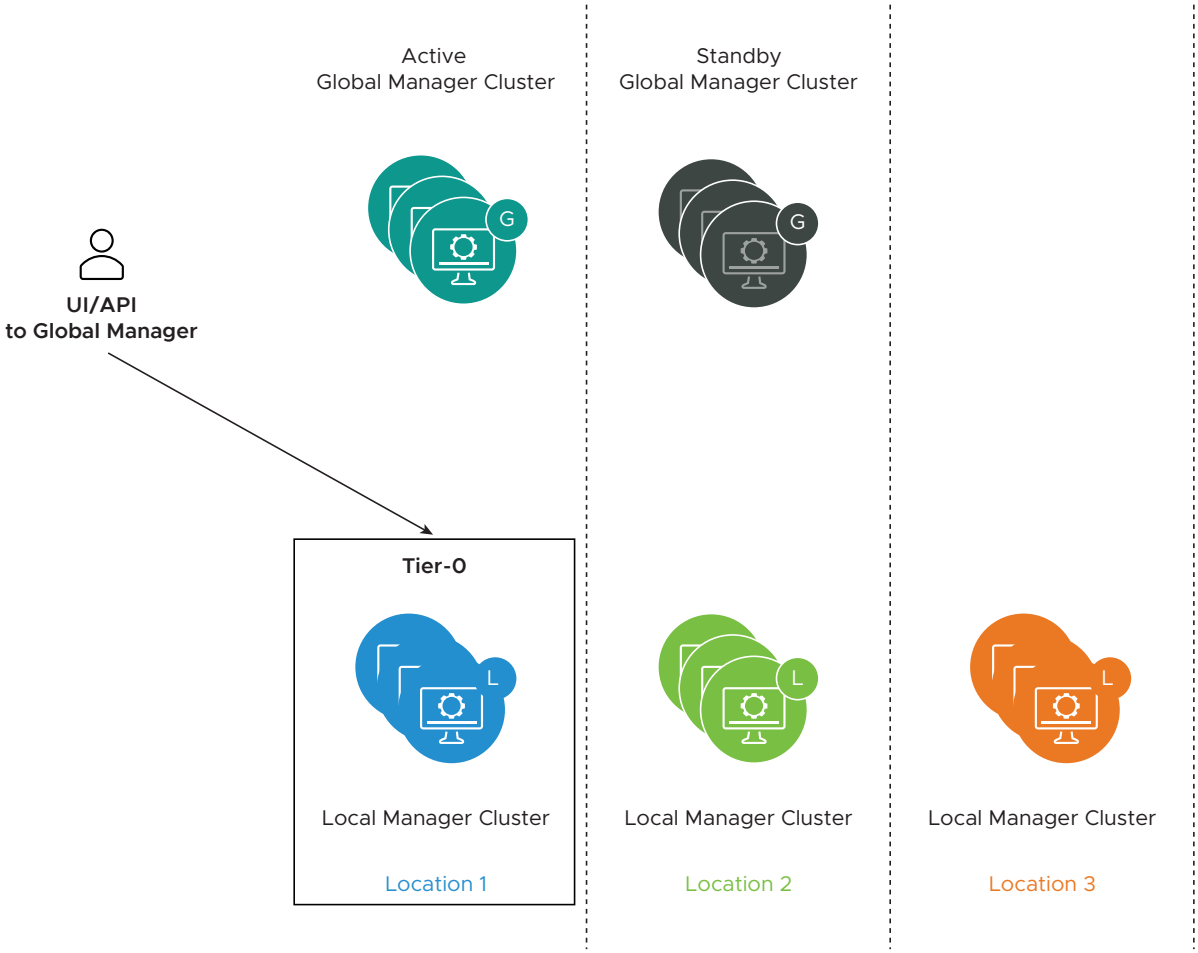


If the tier-0 gateway is added only to Location 1 and Location 2, the configuration is not synced with Location 3.



Making Changes on Local Managers

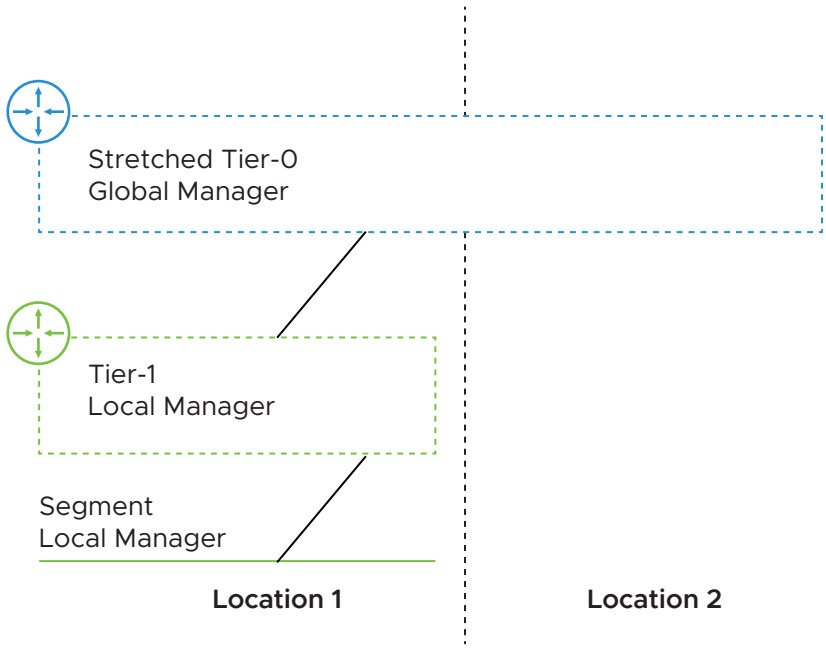
To create objects on a specific Local Manager, you use that Local Manager. These objects do not sync with the active Global Manager or any other Local Managers.



Realizing Global Manager Changes on Local Managers

The Global Manager validates change against the Global Manager and the Local Manager configurations. When a Local Manager receives a configuration from the Global Manager, it realizes the configuration in the fabric nodes of that Local Manager. During this realization, errors or conflicts might get detected. To monitor configuration flow, use the NSX Federation monitoring dashboard. See Monitoring NSX Federation Locations for details.

For example, you can create a tier-0 gateway from Global Manager, and then from a Local Manager you can create and attach a tier-1 gateway to the tier-0 gateway.



Because Local Managers now sync their configurations to the Global Manager, the Global Manager context the tier-0 gateway now appears to be connected. You can delete the tier-0 gateway from the Global Manager, and this change gets synchronized to the Local Managers. When the changes i each location get realized, the following occurs:

- The tier-0 gateway might get deleted from the Local Manager in Location 2.
- The tier-0 gateway might get deleted from the Local Manager in Location 1.
- The tier-0 gateway gets marked for deletion on the Global Manager.

When the tier-0 disconnects from the tier-1 in Location 1, the tier-0 gets deleted from Global Manager.

Most problems are displayed on the user interface. You can also display problems using these API calls.

- On Global Manager:

```
GET /global-manager/api/v1/global-infra/realized-state/alarms
```

- On Local Manager:

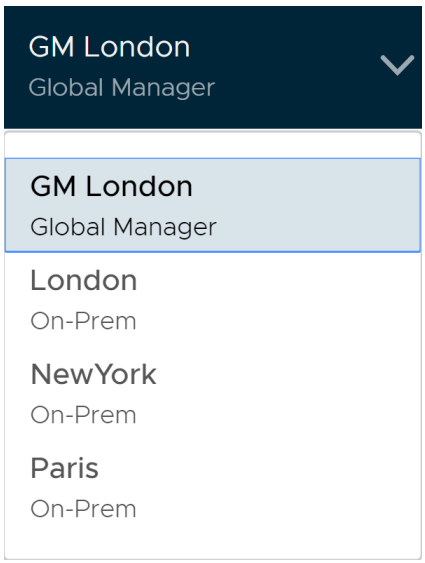
```
GET /policy/api/v1/infra/realized-state/alarms
```

Using the Global and Local Manager Web Interfaces

You can use the Global Manager to create objects that are limited to one location, or span multiple locations. You can also monitor communication details between Local Manager and Global Manager using the Location Manager.

Location Drop-Down Menu on Global Manager

When you log into the active Global Manager web interface, you see a Location drop-down menu in the top navigation bar. Using this menu, you can switch between the active Global Manager and any associated Local Managers. You can also select **System > Location Manager** to view communication activity.

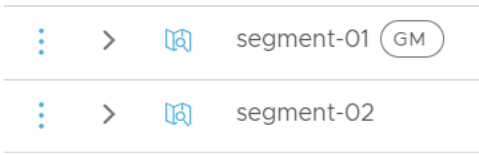


Local and Global Objects

Objects created on a Local Manager are local objects. Objects created from the Global Manager are global objects, though their span might not include all available locations. You can import objects to the Global Manager using the Location Manager.

On a Local Manager, you can see local objects, and any global objects that apply to that location. The global objects have an icon next to them: GM.

This screenshot from the Local Manager web interface shows two segments. The segment `segment-01` has the GM icon next to it, which indicates that it was created on the Global Manager. The segment `segment-02` has no GM icon, which indicates that it was created on the Local Manager.

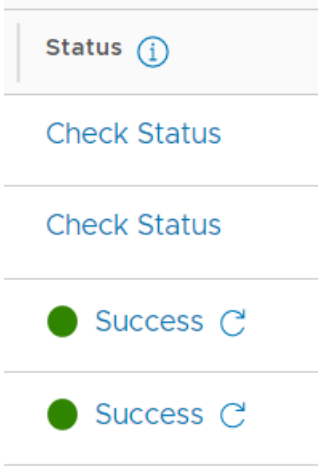


Because all objects on the Global Manager are global, there is no icon displayed when you are logged into the active Global Manager.

Status of Local and Global Objects

Local Managers display the status of both global and local objects. The active Global Manager displays any objects imported from the Local Managers. For details on importing objects from Local to Global Managers using Location Manager, see Monitoring NSX Federation Locations.

To retrieve the latest status from the Local Managers, click **Check Status** for the object. To refresh the status, click the **Refresh** icon.



Monitoring NSX Federation Locations

To monitor the configuration flow between all Global Managers (GM) and Local Managers (LM) across a single NSX Federation federated domain for on-premises solutions you can use the **System > Location Manager** menu or the Location drop-down menu on the GM page. After you configure Locations, alarms display any communication issues between the GM and the LM on the Location Manager page.

You must already have a Location added. See "Add a Location" in the *NSX Installation Guide*.

Use the Location Manager to view latency data, monitor communication synchronizations, check policies, and get better visibility into the health between the different components of the federated domain. The Global Manager View of Location Manager figure shows the Global and Local Manager details. For callout details, see the Callouts for Global Manager View of Location Manager table.

Figure 18-1. Global Manager View of Location Manager

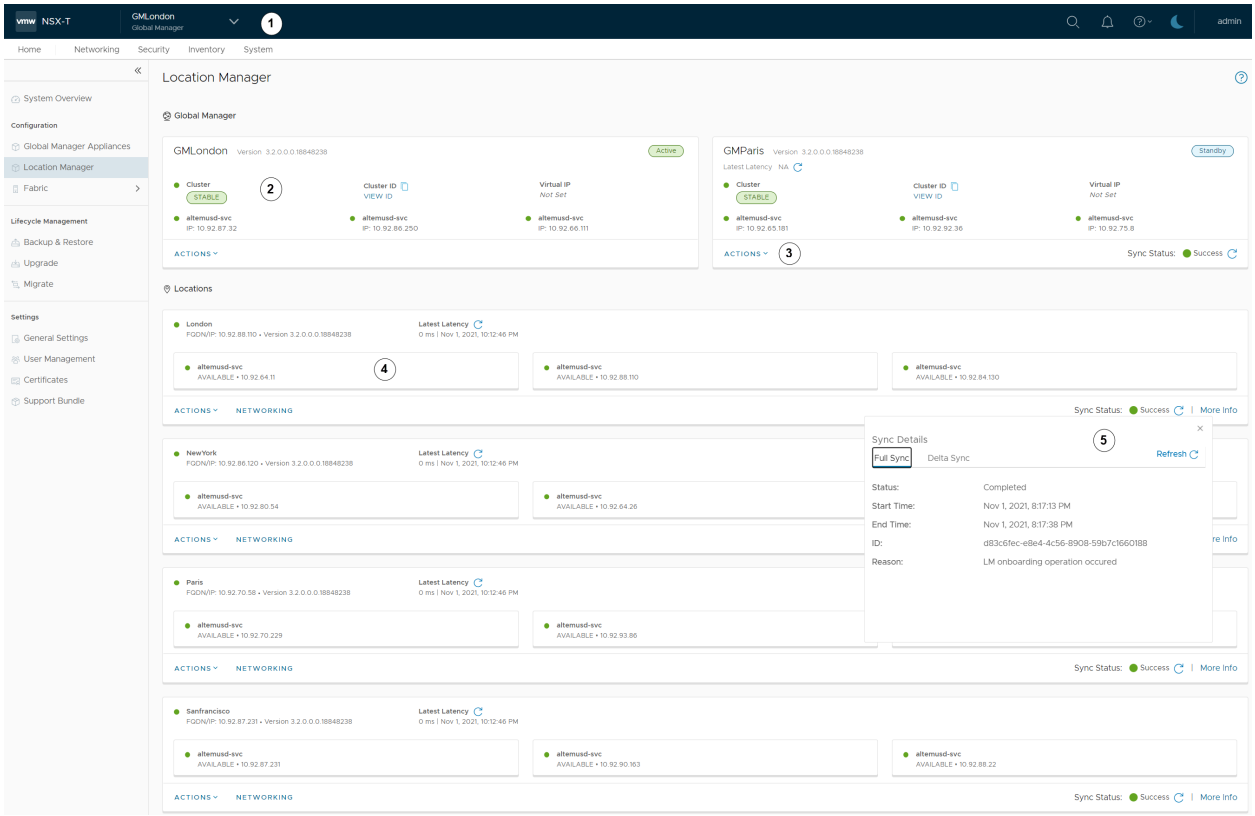


Table 18-3. Callouts for Global Manager View of Location Manager

ID	Options	Description
1	Set location view.	Switch to different GM-active and LM views by selecting dropdown. After logging into the active GM, this drop-down indicates what instance you are logged into and lets you switch to other LM locations views. You can also switch to the standby GM, not using this dropdown, but by selecting the **Actions** menu in the Standby pane and clicking **Access Standby GM**.
2	Active GM information.	■ GM-Active details. ■ Tasks such as GM Active Name.
3	GM Standby information.	■ GM Standby details such as Sync Status. ■ Other details such as GM Standby Name, Access Standby GM View, and Remove GM Standby.

Table 18-3. Callouts for Global Manager View of Location Manager (continued)

ID	Options	Description
4	LM pane information.	■ LM details such as Latency Data, Sync Status, and Networking RTEP configuration. ■ Tasks such as edit settings, remove, evacuate location, and import to GM. Evacuate location allows network recovery for sites that have not lost communication. For example, if you want to migrate a data center or perform a disaster recovery test. Import to GM allows the former LM configuration to be pushed up to GM (see "Importing Configurations from Local Manager" in the *NSX Installation Guide*.
5	LM More Info pane.	■ View detailed Full Sync from GM Active to LM information such as Start Time, End Time, and Reason. ■ View detailed Delta Sync from GM Active to LM information such as Message Queues on GM and LM.

The Local Manager View of Location Manager figure shows an example of the window details. For callout details, see the Callouts for Local Manager View table.

Figure 18-2. Local Manager View of Location Manager

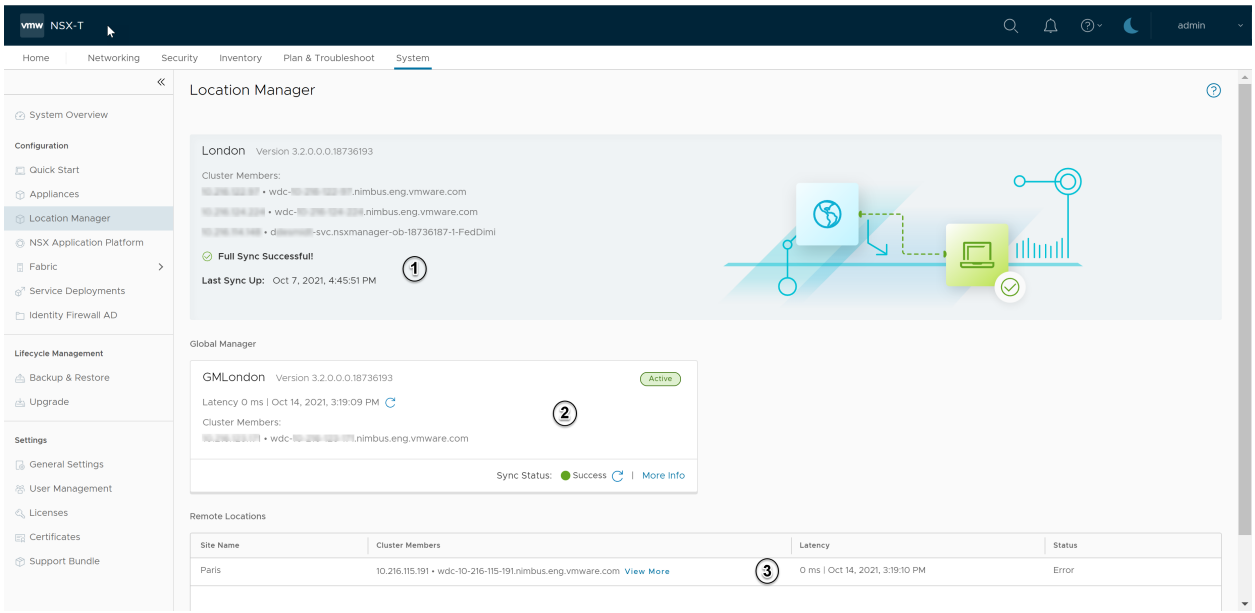


Table 18-4. Callouts for Local Manager View of Location Manager

ID	Task	Description
1	View LM details.	NSX version, cluster IPs and names, status of full sync and when it completed.
2	View related active GM details.	■ Software versions, latency details between GM and LM, cluster details, and color-coded sync status. ■ Green is communication is OK. Orange is some backlog in the queue based on threshold. Red is sync is not occurring. ■ To see the full sync details including status, start and end times, ID, and status reasons, click More Info.
3	View remote location details.	Cluster details, including View More information, latency data between LM and GM, and current status.

The Supported Communication Channel Details table provides details about the Location Manager displays.

Table 18-5. Supported Communication Channel Details

Supported Communication Channels	Display Details

All channels	■ Status (color-coded)
	  ■ Green - OK.
	  ■ Orange - Sync between active GM and Standby GM is in progress.
	  ■ Red - Communication is lost.
	■ When latency is over 500 ms, an alarm gets raised. For NSX Federation events, see the Event Catalog.

On a GM and LM UI for each remote LM	■ Current queue levels contain the number of queued messages compared to the queue size. ■ Last measured latency - Uses ICMP or TCP traffic between the local LM and the remote LM to measure the network round trip in milliseconds (ms).
On a LM UI for a GM active cluster and a GM standby cluster	■ Current queue levels contain the number of queued messages compared to the queue size for GM active only. ■ Last measured latency - Measures the network round trip in ms using ICMP or TCP traffic between the local LM and the active and standby GM. ■ When latency is over 500 ms, an alarm gets raised.
On an active GM UI for a standby GM cluster	■ When Active-Standby GM replication is not working, an alarm gets raised.

Overriding Global Manager Configurations on Local Manager

When you create an object from Global Manager, the same configuration is propagated to all relevant locations. You can override some Global Manager configurations on a Local Manager.

To override a configuration, click the three dots menu ( ) next to the configuration, and click **Edit**. If the **Edit** menu item is dimmed, you cannot override this configuration.

If a configuration is overridden, you see this icon in the status column on both Global Manager and Local Manager: .

To remove an override, click the three dots menu ( ) next to the configuration, and click **Revert**. The configuration from Global Manager is restored.

If you override a configuration from Global Manager on a Local Manager, and then you delete the configuration from the Global Manager, the configuration persists on the Local Manager. When you revert the configuration, the configuration is deleted from Local Manager.

You can get a list of all configurations that have been overridden. Make this API request to the Global Manager: `GET https://<global-mgr>/global-manager/api/v1/global-infra/overridden-resources`.

Gateway Configurations

Gateway configurations are found in **Networking > Tier-0 Gateways** and **Networking > Tier-1 Gateways**.

You can modify the following gateway configurations:

- Tier-0 Gateway BGP Configuration
- Tier-0 Gateway Interfaces

Profile Configurations

Profile configurations on Global Manager are used in all Local Managers. There is no span setting for a profile configuration.

You can override the following global profile configurations from Local Manager:

- Segment Profiles: **Networking > Segments > Segment Profiles**
  - IP Discovery Profiles
  - MAC Discovery Profiles
  - Segment Security Profiles
  - SpoofGuard Profiles
- Networking Profiles: **Networking > Networking Settings**
  - IPv6 DAD Profiles
  - IPv6 ND Profiles
  - Gateway QoS Profiles
  - BFD Profiles
- Context Profiles: **Inventory > Context Profiles**

- Security Profiles: **Security > Security Profiles**
  - Firewall Session Timer Profile
  - Edge Gateway Flood Protection Profiles
  - Firewall Flood Protection Profiles
  - DNS Security Profiles
  - CPU and Memory Threshold Profiles are API only:
    - Override with `PUT/PATCH https://<local-manager>/policy/api/v1/global-infra/settings/firewall/cpu-mem-thresholds-profiles/<id>?action=override`.
    - Revert with `DELETE https://<local-manager>/policy/api/v1/global-infra/settings/firewall/cpu-mem-thresholds-profiles/<id>`.
- Troubleshooting Profiles: **Plan & Troubleshoot**
  - Firewall IPFIX Profiles
  - Switch IPFIX Profiles
  - IPFIX Firewall Collector
  - IPFIX Switch Collector
  - Remote L3 Span Port Mirroring Profile
  - Logical Span Port Mirroring Profile
  - QoS Profile

Networking in NSX Federation

Tier-0 gateways, tier-1 gateways, and segments can span one or more locations in the NSX Federation environment.

When you plan your network topology, keep these requirements in mind:

- Tier-0 and tier-1 gateways can have a span of one or more locations.
- The span of a tier-1 gateway must be equal to, or a subset of, the span of the tier-0 gateway it is attached to.
- A segment has the same span as the tier-0 or tier-1 gateway it is attached to. Isolated segments are not realized until they are connected to a gateway.
- NSX Edge nodes in the Edge Cluster selected on the Global Manager for tier-0 and tier-1 gateways must be configured with the Default TZ Overlay.

You can create different topologies to achieve different goals.

- You can create segments and gateways that are specific to a given location. Each site has its own configuration, but you can manage everything from the Global Manager interface.

- You can create segments and gateways that span locations. These stretched networks provide consistent networking across sites.

Tier-0 Gateway Configurations in NSX Federation

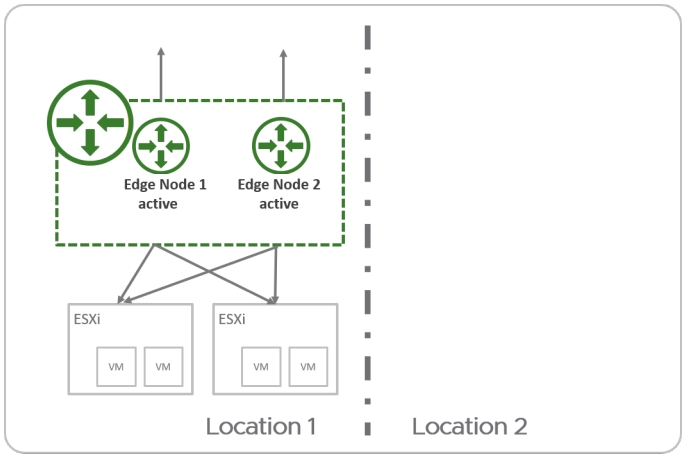
With NSX Federation, you can deploy a tier-0 gateway that is limited to a single location, or you can stretch it across multiple locations.

Tier-0 gateways can have one of the following configurations:

- Non-stretched tier-0 gateway.
- Stretched active-active with primary and secondary locations.
- Stretched active-active with all primary locations.
- Stretched active-standby with primary and secondary locations.

Non-Stretched Tier-0 Gateway

You can create a tier-0 gateway from Global Manager that spans only one location. This is similar to creating the tier-0 gateway on the Local Manager directly, but has the advantage that you can manage it from Global Manager.



Stretched Active-Active Tier-0 Gateway with Primary and Secondary Locations

In an active-active tier-0 gateway with primary and secondary locations, the following applies:

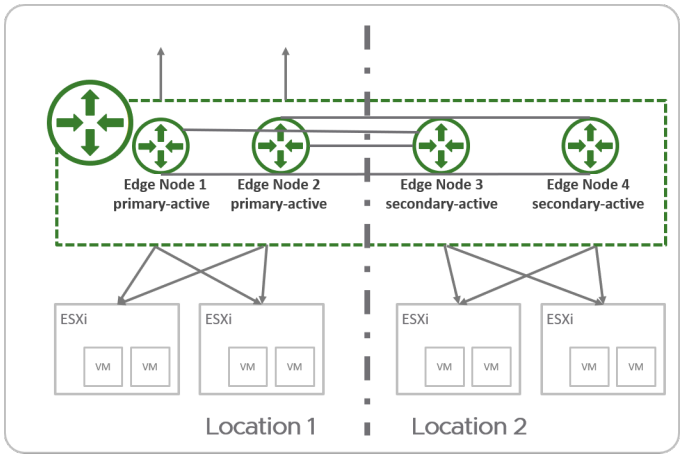
- All Edge nodes are active at the same time, therefore the tier-0 cannot run stateful services.
- All traffic enters and leaves through the Edge nodes in the primary location.

If both the tier-0 gateway and the linked tier-1 gateway have primary and secondary locations, configure the same location to be primary for both gateways to reduce cross-location traffic.

Important In this topology, NSX ensures that all egress traffic leaves through the primary location.

If your environment has stateful services, such as external firewall, on the physical network, you must ensure that the return traffic enters through the primary location. For example, you can add AS path prepending on the BGP peers in your secondary locations.

If you do not have stateful services on your physical network, and you choose to have asymmetric routing, you must disable Unicast Reverse Path Forwarding (uRPF) on all externally tier-0 interfaces.

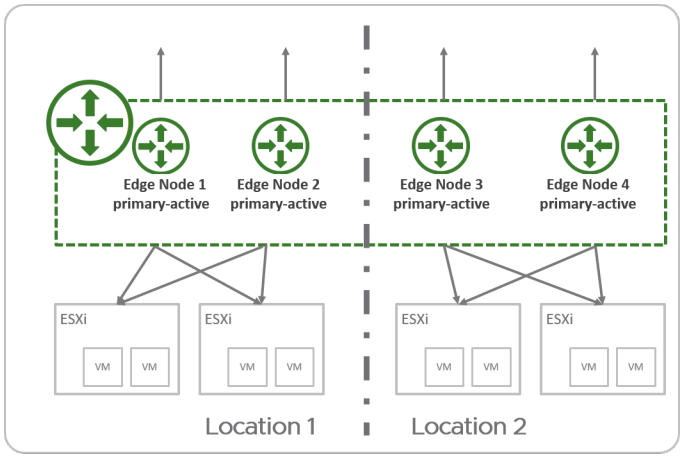


Stretched Active-Active Tier-0 Gateway with All Primary Locations

In an active-active tier-0 gateway with all primary locations, the following applies:

- All Edge nodes are active at the same time, therefore the tier-0 cannot run stateful services.
- All traffic enters and leaves through Edge nodes in the same location as the workloads.

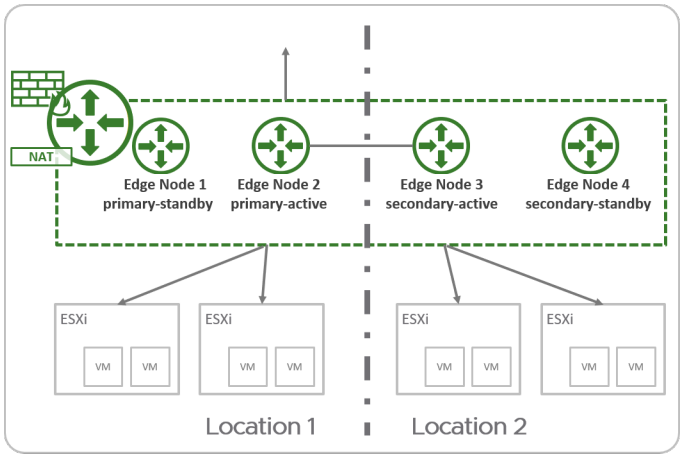
Important This topology allows traffic to egress locally from each location. You must ensure that return traffic enters the same location to allow stateful services such as firewall. For example, you can configure a location-specific NAT IP so that return traffic is always routed back to the same location that it left.



Stretched Active-Standby Tier-0 Gateway with Primary and Secondary Locations

In an active-standby tier-0 gateway with primary and secondary locations, the following applies:

- Only one Edge node is active at a time, therefore the tier-0 can run stateful services.
- All traffic enters and leaves through the active Edge node in the primary location.



For Active Standby tier-0 gateways, the following services are supported:

- Network Address Translation (NAT)
- Gateway Firewall
- DNS
- DHCP

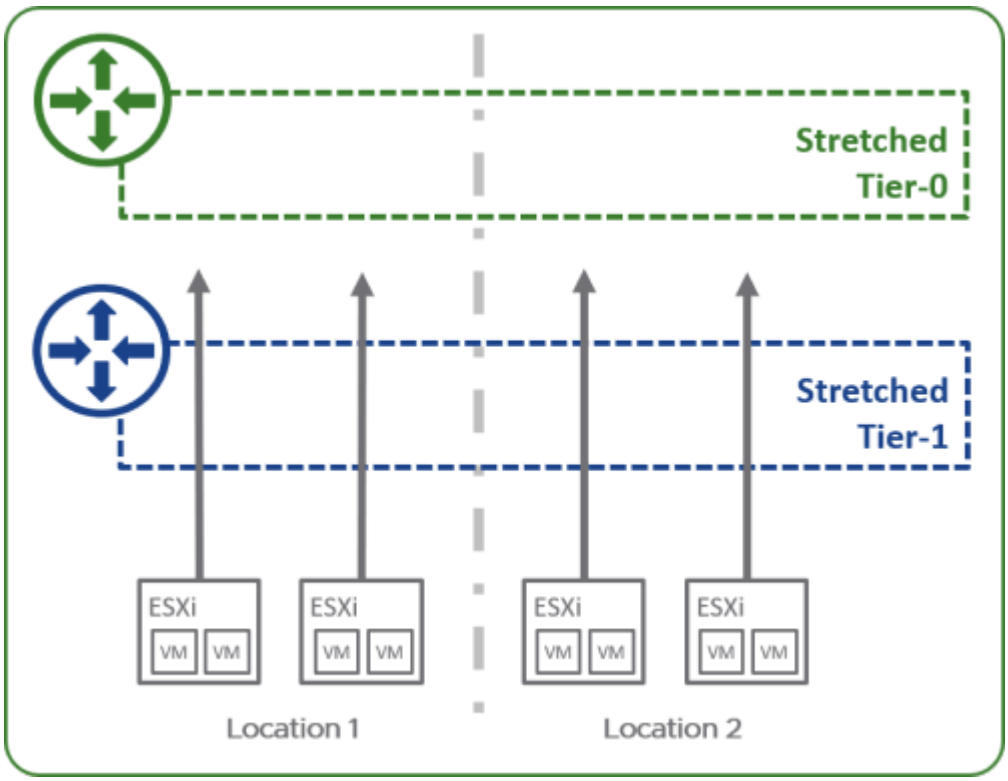
See Features and Configurations Supported in NSX Federation for more information.

Tier-1 Gateway Configurations in NSX Federation

With NSX Federation, you can deploy a tier-1 gateway to provide distributed routing only, or you can configure services on it.

Tier-1 Gateway for Distributed Routing Only

You can create a tier-1 gateway in NSX Federation for distributed routing only. This gateway has the same span as the tier-0 gateway it is linked to. The tier-1 does not use Edge nodes for routing. All traffic is routed from host transport nodes to the tier-0 gateway. However, to enable cross-location forwarding, the tier-1 allocates two Edge nodes from the Edge cluster configured on the linked tier-0 to use for that traffic.



Tier-1 Gateway with Services or Custom Span

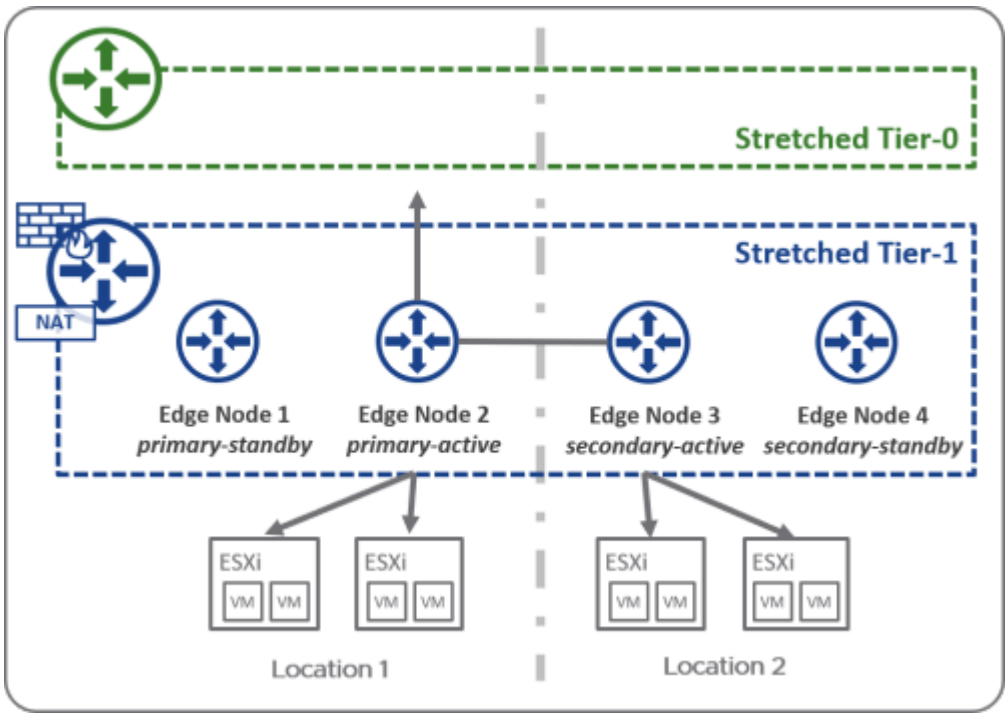
You configure the tier-1 gateway with Edge clusters if you need one of the following configurations:

- You want to run services on the tier-1 gateway.
- You want to deploy a tier-1 gateway that has a different span than the linked tier-0 gateway.

  You can remove locations, but you cannot add locations that are not already included the span of the tier-0 gateway.

You select one of the locations to be the primary location. All other locations are secondary. The HA mode for the tier-1 gateway is Active Standby. All traffic passing through this tier-1 gateway passes through the active edge node in the primary location.

If both the tier-1 gateway and the linked tier-0 gateway have primary and secondary locations, configure the same location to be primary for both gateways to reduce cross-location traffic.



Configure Edge Nodes for Stretched Networking

If you want to create gateways and segments that span more than one location, you must configure a remote tunnel endpoint (RTEP) on Edge nodes in each location.

When you configure an RTEP, do it on an Edge cluster basis. All Edge nodes in the cluster must have an RTEP configured. You do not need to configure all Edge clusters with RTEP. RTEPs are required only if the Edge cluster is used to configure a gateway that spans more than one location.

You can configure the TEP and RTEP to use the same physical NIC on the Edge node or use separate physical NICs.

This procedure describes this task starting from your Local Manager. You can also configure RTEPs from your Global Manager by using the Location Manager site selection drop-down to choose the Local Manager.

Prerequisites

- Verify that each location participating in the stretched network has at least one Edge cluster.
- Determine which layer 3 networks and VLANs to use for RTEP networks.
  - Intra-location tunnel endpoints (TEP) and inter-location tunnel endpoints (RTEP) must use separate VLANs and layer 3 subnets.
- Verify that all RTEP networks used in a given NSX Federation environment have IP connectivity to each other.
- Verify that external firewalls allow cross-location RTEP tunnels. See VMware Ports and Protocols at https://ports.vmware.com/home/NSX.

- Configure the MTU for RTEP on each Local Manager. The default is 1500. Set the RTEP MTU to be as high as your physical network supports. On each Local Manager, select **System > Fabric > Settings**. Click **Edit** next to **Remote Tunnel Endpoint**.
- Optionally, if you do not use DHCP for RTEP, configure the RTEP IP pool for your site in order to configure the RTEPs for the Edge cluster. For details, go to "Add an IP Address Pool" in the *NSX Administration Guide*.

Procedure

1 From your browser, log in with admin privileges to the Local Manager at https://<local-manager-ip-address>.

2 To configure a new RTEP, select **System > Quick Start**.

3 Click **Configure Remote Tunnel Endpoint > Getting Started**.

4 You can select all Edge Nodes in this cluster or one node at a time. Provide the following details for the RTEP configuration:

Option	Description
Edge Switch	Select an edge switch from the drop-down menu.
Teaming Policy Name	(Optional) Select a teaming policy if you have one configured.
RTEP VLAN	Enter the VLAN ID for the RTEP network. Valid values are between 1 and 4094.
IP Assignment	Select an option from the drop-down. For example, select Use IP Pool and choose an option from the drop-down list.
IP Pool for all nodes	Select an IP pool for all nodes in this Edge Cluster. If you want to assign an IP address to an individual node, you can edit the RTEP configuration later.
Inter Location MTU	The default is 1500.

5 Click **Save**.

The green notification banner shows that all the Edge nodes in this edge cluster have been configured successfully.

6 To add the RTEPs to the edge cluster for your other site locations, repeat steps 2 through 5.

7 To view or edit an existing Edge transport node:

a Select **System > Fabric > Nodes > Edge Transport Nodes**.

b Select an Edge node and click **Tunnels**. If an RTEP is configured, it is displayed in the **Remote Tunnel Endpoint** section.

c Click **Edit** to modify the RTEP configuration.

The Configure Edge Nodes for Stretched Networking screen opens in the Local Manager with that Edge cluster selected.

Add a Tier-0 Gateway from Global Manager

You can add a tier-0 gateway from the Global Manager. This gateway can have a span of one or more locations. This span affects the span of the tier-1 gateways and segments attached to it.

See Tier-0 Gateway Configurations in NSX Federation for details about tier-0 gateway configurations in NSX Federation.

The following settings must be kept consistent across locations. If you change these settings from the Global Manager web interface, those changes are automatically applied on all locations. However, if you change these settings using the API, you must manually make the same changes in each location.

- Local AS
- ECMP settings
- Multipath Relax settings
- Graceful Restart

Important When you create a tier-0 gateway from Global Manager, you must configure an external interface in each location that the tier-0 is stretched to. Each external interface must be connected to a segment that was created from Global Manager, with the **Connectivity** set to None and the **Traffic type** set to VLAN. See Add a Segment from Global Manager. The Edge nodes configured with those external interfaces are used for inter-location communication, even if northbound communication is not needed.

Prerequisites

- If you are creating a tier-0 gateway that spans more than one location, verify that each location has Edge nodes configured with RTEPs for stretched networking. See Configure Edge Nodes for Stretched Networking.

Procedure

1. From your browser, log in with admin privileges to the active Global Manager at https://<global-manager-ip-address>.
2. Select **Networking > Tier-0 Gateways**.
3. Enter a name for the gateway.
4. Select an HA (high availability) mode to configure within each location.

   The default mode is active-active. In the active-active mode, traffic is load balanced across edge nodes in all locations. In the active-standby mode, an elected Edge node processes traffic in each location. If the active node fails, the standby node becomes active.

   Note Active-standby tier-0 gateways are supported starting in NSX 3.0.1.

5 If the HA mode is active-standby, select a failover mode.

Option	Description
Preemptive	If the preferred node fails and recovers, it will preempt its peer and become the active node. The peer will change its state to standby.
Non-preemptive	If the preferred node fails and recovers, it will check if its peer is the active node. If so, the preferred node will not preempt its peer and will be the standby node.

6 Specify the span of this tier-0 gateway by providing the following details for each location. To add additional locations, click **Add Location**.

Option	Description
Location	Select the location from the drop-down menu.
Edge Cluster	Select an Edge cluster from this location. If you are configuring a stretched tier-0, you must select an Edge cluster that contains Edge nodes that are configured with an RTEP.
Mode	Each location of the tier-0 gateway can have a mode of **Primary** or **Secondary**. ■ If the HA mode is **Active Active**, you can configure the tier-0 gateway with all locations mode set to primary. 1 Select the **Mark all locations as Primary** toggle to mark all locations as primary. ■ If the HA mode is **Active Active** or **Active Standby**, you can configure the tier-0 gateway with one location set to **Primary**, and all others set to **Secondary**. 1 Select **Primary** mode for one location. In all other locations, set mode to **Secondary**. 2 For secondary locations, you must select a fallback preference.

7 Click **Additional Settings**.

a In the **Internal Transit Subnet** field, enter a subnet.

This is the subnet used for communication between components within this gateway. The default is 169.254.0.0/24.

b In the **T0-T1 Transit Subnets** field, enter one or more subnets.

These subnets are used for communication between this gateway and all tier-1 gateways that are linked to it. After you create this gateway and link a tier-1 gateway to it, you will see the actual IP address assigned to the link on the tier-0 gateway side and on the tier-1 gateway side. The address is displayed in **Additional Settings > Router Links** on the tier-0 gateway page and the tier-1 gateway page. The default is 100.64.0.0/16.

After the tier-0 gateway is created, you can change the **T0-T1 Transit Subnets** by editing the gateway. Note that this will cause a brief disruption in traffic.

c In the **Intersite Transit Subnet** field, enter a subnet. This subnet is used for cross-location communication between gateway components. The default is 169.254.32.0/20.

8 Click **Save**.

9 To configure interfaces, click **Interfaces** and **Set**. Configure an external interface for each location that the tier-0 gateway spans.

a Click **Add Interface**.

b Enter a name.

c Select a location.

d Select a type.

If the HA mode is active-standby, the choices are **External**, **Service**, and **Loopback**. If the HA mode is active-active, the choices are **External** and **Loopback**.

Service interfaces are supported only on gateways that span one location. If the gateway is stretched, service interfaces are not supported.

e Enter an IP address in CIDR format.

f Select a segment.

The segment must be created from the Global Manager, with the **Connectivity** set to None and the **Traffic type** set to VLAN. See Add a Segment from Global Manager.

g If the interface type is not **Service**, select an NSX Edge node.

h (Optional) If the interface type is not **Loopback**, enter an MTU value.

i Skip **PIM** configuration.

Multicast is not supported in NSX Federation.

j (Optional) Add tags and select an ND profile.

k (Optional) If the interface type is **External**, for **URPF Mode**, you can select **Strict** or **None**.

URPF (Unicast Reverse Path Forwarding) is a security feature.

l (Optional) After you create an interface, you can download the aggregate of ARP proxies for the gateway by clicking the menu icon (three dots) for the interface and selecting **Download ARP Proxies**.

You can also download the ARP proxy for a specific interface by expanding a gateway and then expanding **Interfaces**. Click an interface and click the menu icon (three dots) and select **Download ARP Proxy**.

Note You cannot download the ARP proxy for loopback interfaces.

10 Click **Routing** to add IP prefix lists, community lists, static routes, and route maps.

When you add a static route on a tier-0 gateway, the default behavior is that the static routes are pushed to all locations configured on the gateway. However, the routes are enabled only on the primary locations. This ensures that on the secondary locations, the routes that are learned from the primary location are preferred.

If you want to change this behavior, you can use the **Enabled on Secondary** setting and the **Scope** setting.

If you select **Enabled on Secondary**, the static route is also enabled on the secondary locations.

When you add a next hop for a static route, you can set the **Scope**. The scope can be an interface, a gateway, or a segment. On a tier-0 gateway created from Global Manager, the scope can also be a location. You can use the scope setting to configure different next hops for each location.

11 Click **BGP** to configure BGP.

When you configure BGP on a tier-0 gateway from the Global Manager, most settings apply to all locations.

Some of the settings within the BGP configuration, such as **Route Aggregation** and **BGP Neighbors** prompt you to provide separate values for each location.

See Configure BGP for more information about configuring BGP.

12 To configure route redistribution, click **Route Redistribution**, and for each location, click **Set**.

Select one or more of the sources:

- Tier-0 subnets: **Static Routes**, **NAT IP**, **IPSec Local IP**, **DNS Forwarder IP**, **EVPN TEP IP**, **Connected Interfaces & Segments**.

  Under **Connected Interfaces & Segments**, you can select one or more of the following: **Service Interface Subnet**, **External Interface Subnet**, **Loopback Interface Subnet**, **Connected Segment**.

- Advertised tier-1 subnets: **DNS Forwarder IP**, **Static Routes**, **LB VIP**, **NAT IP**, **LB SNAT IP**, **IPSec Local Endpoint**, **Connected Interfaces & Segments**.

  Under **Connected Interfaces & Segments**, you can select **Service Interface Subnet** and/or **Connected Segment**.

What to do next

Set up a tier-1 gateway from Global Manager.

Add a Tier-1 Gateway from Global Manager

A gateway can be configured in one or more locations. These locations are the span of the gateway. A tier-1 gateway cannot have a greater span than the tier-0 gateway it is connected to.

See Tier-1 Gateway Configurations in NSX Federation for details about tier-1 gateway configuration options in NSX Federation.

Prerequisites

Verify you have a tier-0 gateway configured.

Procedure

1 From your browser, log in with admin privileges to an NSX Manager at https://<global-manager-ip-address>.

2 Select **Networking > Tier-1 Gateways**.

3 Click **Add Tier-1 Gateway**.

4 Enter a name for the gateway.

5 Select a tier-0 gateway to connect to this tier-1 gateway to create a multi-tier topology.

- If you select a tier-0 gateway, the Locations configuration is populated with the same locations that are configured on the tier-0. If needed, you can modify the locations configuration in the Locations section.
- If you do not select a tier-0 gateway, you can select locations. However, if you later connect the tier-1 gateway to a tier-0 gateway, you might need to update the locations to create a valid configuration.

6 In **Locations**, you can change the **Enable Edge Clusters for Services or Custom Span** setting. It is disabled by default.

- Leave **Enable Edge Clusters for Services or Custom Span** disabled if you want the tier-1 gateway to have the same span as the tier-0 gateway, and you do not need to enable services on the tier-1 gateway. The tier-1 gateway will perform distributed routing only.
- Enable **Enable Edge Clusters for Services or Custom Span** if you want to choose a subset of locations for the tier-1 gateway, or if you want to enable services on the tier-1 gateway.

If you enable **Enable Edge Clusters for Services or Custom Span**, enter the location, cluster, and mode information.

a Select a location from the drop-down menu. If you linked this tier-1 gateway to a tier-0 gateway, the locations of that tier-0 gateway are automatically listed. If needed, you can delete a location.

b Select an NSX Edge cluster for each location. If the tier-1 gateway spans more than one location, the Edge clusters must already be configured with an RTEP for each of its Edge Nodes.

c (Optional) To select specific Edge nodes, click **Set** next the Edge cluster.

Edge nodes are automatically allocated if you do not select Edge nodes.

d Select a mode for each location. Mode can be Primary or Secondary.

Only one location can be configured with Primary mode. All northbound traffic from this tier-1 gateway is sent through this location.

7 If you have enabled Edge clusters, select a failover mode.

Option	Description
Preemptive	If the preferred NSX Edge node fails and recovers, it will preempt its peer and become the active node. The peer will change its state to standby.
Non-preemptive	If the preferred NSX Edge node fails and recovers, it will check if its peer is the active node. If so, the preferred node will not preempt its peer and will be the standby node. This is the default option.

8 Skip selecting a size from the **Edge Pool Allocation Size** drop-down menu.

9 If you have enabled Edge clusters, select a setting for **Enable StandBy Relocation**.

Standby relocation means that if the Edge node where the active or standby logical router is running fails, a new standby logical router is created on another Edge node to maintain high availability. If the Edge node that fails is running the active logical router, the original standby logical router becomes the active logical router and a new standby logical router is created. If the Edge node that fails is running the standby logical router, the new standby logical router replaces it.

10 (Optional) Click **Route Advertisement**.

Select one or more of the following:

- **All Static Routes**
- **All NAT IP's**
- **All DNS Forwarder Routes**
- **All LB VIP Routes**
- **All Connected Segments and Service Ports**
- **All LB SNAT IP Routes**
- **All IPSec Local Endpoints**

11 Click **Save**.

12 (Optional) Click **Route Advertisement**.

a In the **Set Route Advertisement Rules** field, click **Set** to add route advertisement rules.

13 (Optional) Click **Additional Settings**.

a For IPv6, you can select or create an **ND Profile** and a **DAD Profile**.

These profiles are used to configure Stateless Address Autoconfiguration (SLAAC) and Duplicate Address Detection (DAD) for IPv6 addresses.

b Select an **Ingress QoS Profile** and an **Egress QoS Profile** for traffic limitations.

These profiles are used to set information rate and burst size for permitted traffic. See Add a Gateway QoS Profile for more information on creating QoS profiles.

If this gateway is linked to a tier-0 gateway, the **Router Links** field shows the link addresses.

14 (Optional) Click **Service Interfaces** and **Set** to configure connections to segments. Required in some topologies such as VLAN-backed segments or one-arm load balancing.

   Service interfaces are supported only on gateways that span one location. If the gateway is stretched, service interfaces are not supported.

   a Click **Add Interface**.

   b Enter a name and IP address in CIDR format.

      If you configure multicast on this gateway, you must not configure tier-1 addresses as static RP address in the PIM profile.

   c Select a segment.

   d In the **MTU** field, enter a value between 64 and 9000.

   e For **URPF Mode**, you can select **Strict** or **None**.

      URPF (Unicast Reverse Path Forwarding) is a security feature.

   f Add one or more tags.

   g In the **ND Profile** field, select or create a profile.

   h Click **Save**.

   i (Optional) After you create an interface, you can download the ARP proxies for the gateway by clicking the menu icon (three dots) for the interface and selecting **Download ARP Proxies**.

      You can also download the ARP proxy for a specific interface by expanding a gateway and then expanding **Service Interfaces**. Click an interface and click the menu icon (three dots) and select **Download ARP Proxy**.

15 (Optional) Click **Static Routes** and **Set** to configure static routes.

   a Click **Add Static Route**.

   b Enter a name and a network address in the CIDR or IPv6 CIDR format.

   c Click **Set Next Hops** to add next hop information.

   d Click **Save**.

Add a Segment from Global Manager

You can add two kinds of segments: overlay-backed segments and VLAN-backed segments. When you create segments from Global Manager, only overlay-backed segments can span multiple locations.

You can view segments ports from Global Manager, but you cannot create or modify them. If you need to create or modify a segment port, you must do it from the Local Manager.

Important Do not change the gateway connectivity of a segment in NSX Federation. Changing the gateway affects the span of the segment. If the span changes in such a way that it excludes a location, the segment is deleted on the excluded location. You must disconnect all VMs before you shrink the span of a segment.

Prerequisites

Verify that each location has a default overlay transport zone configured. The default overlay transport zone is used to create global overlay segments. From each Local Manager, select **System > Fabric > Transport Zones**. Select an overlay transport zone, and click **Actions > Set as Default Transport Zone**.

Procedure

1. From your browser, log in with admin privileges to a Global Manager at https://<global-manager-ip-address>.
2. Select **Networking > Segments**.
3. Click **Add Segment**.
4. Enter a name for the segment.
5. Select the Connectivity, Traffic Type, and Locations for this segment.

Table 18-6. Segment Configurations

Connectivity	Traffic Type	Location and Transport Zone	Details
A global tier-0 or tier-1 gateway	Overlay	The **Location** section is populated with the following configurations: ■ the same locations that are configured on the attached gateway. ■ the default overlay transport zone for each location.	Use this configuration to create a global overlay-backed segment connected to the selected global gateway.
None	VLAN	You must select one location for this segment. You must also select a transport zone from that location.	Use this configuration to create a global VLAN-backed segment to use for a tier-0 external interface.
None	Overlay	No locations or transport zones can be selected.	This segment is created on the Global Manager but is not realized in any Local Managers. You can attach it to a gateway later.

Creating a VLAN-backed segment that is attached to a gateway is not supported.

6 Enter the Gateway IP address of the subnet in a CIDR format. A segment can contain an IPv4 subnet, or an IPv6 subnet, or both.

- If a segment is not connected to a gateway, subnet is optional.
- If a segment is connected either to a tier-1 or tier-0 gateway, subnet is required.

Subnets of one segment must not overlap with the subnets of other segments in your network. A segment is always associated with a single virtual network identifier (VNI) regardless of whether it is configured with one subnet, two subnets, or no subnet.

7 Skip **Set DHCP Config**.

Only static bindings are supported on a segment created from Global Manager. See Features and Configurations Supported in NSX Federation.

8 If the transport zone is of type VLAN, specify a list of VLAN IDs. If the transport zone is of type Overlay, and you want to support layer 2 bridging or guest VLAN tagging, specify a list of VLAN IDs or VLAN ranges

9 (Optional) Select an uplink teaming policy for the segment.

This drop-down menu displays the named teaming policies, if you have added them in the VLAN transport zone. If no uplink teaming policy is selected, the default teaming policy is used.

- Named teaming policies are not applicable to overlay segments. Overlay segments always follow the default teaming policy.
- For VLAN-backed segments, you have the flexibility to override the default teaming policy with a selected named teaming policy. This capability is provided so that you can steer the infrastructure traffic from the host to specific VLAN segments in the VLAN transport zone. Before adding the VLAN segment, ensure that the named teaming policy names are added in the VLAN transport zone.

10 Click **Save**.

11 To continue configuring the segment, click **Yes** when prompted.

12 To select segment profiles, click **Segment Profiles**.

13 Click **Save**.

Configure Bridging on Global Manager

Bridging allows the communication between overlay segments and physical VLANs. Starting with NSX 3.2.2, you can configure edge bridges from the Global Manager and apply them to stretched and non-stretched segments.

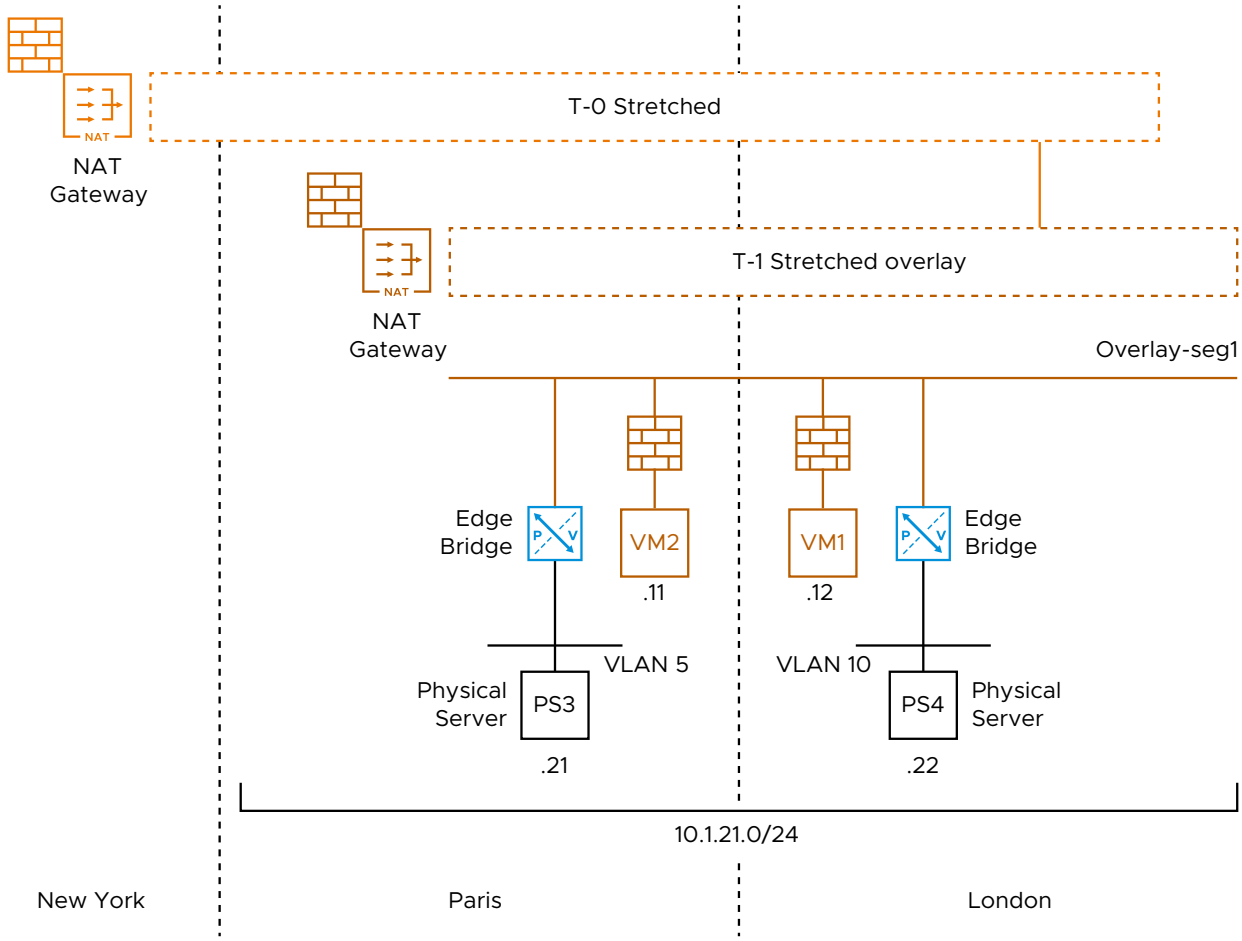
For example, in the logical diagram, Edge Bridging in NSX:

- VM2 connects to an overlay segment (overlay-seg1) and a physical server (PS3) on physical VLAN5. All are in the subnet 10.1.21.0/24.

- VM1 connects to an overlay segment (overlay-seg1) and physical server (PS4) on physical VLAN10. All are on the subnet 10.1.21.0/24.

The edge bridge allows the communication between overlay and VLAN.

Figure 18-3. Edge Bridging in NSX



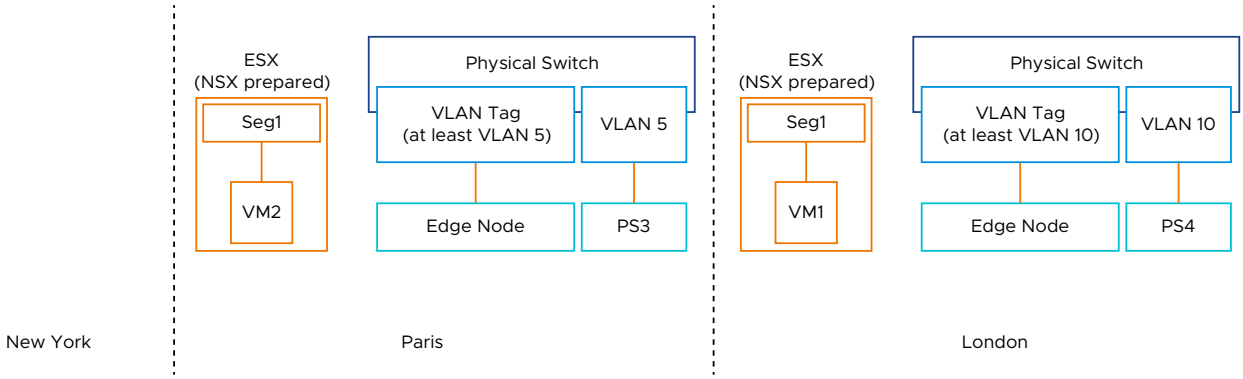
This functionality includes the following support as depicted in the edge bridging in NSX Federation diagram:

- Bridge creation on a segment on a given location.
- The segment can span one site or multiple sites. For instance, if you had three sites (New York, Paris, and London), you have a bridge segment, overlay-seg1, that spans Paris and London.
- The bridge is present on one site in one edge cluster with an edge node on the same site. For instance, the bridging profile bridge-paris defines the bridge in Paris.
- Multiple bridges can be assigned to the same segment. For example, you can bridge the segment overlay-seg1 in Paris with bridge-Paris to a given VLAN, such as VLAN 5.

You can use overlay/VLAN bridging on edge nodes which you configure using edge bridge profiles. The edge bridge profile contains the edge node primary and edge node backup of an edge cluster. You then associate that edge bridge profile in your overlay-segment.

Note As shown in the NSX Edge Bridge Physical View diagram, the edge nodes must have connectivity to those physical VLANs (VLAN5 and VLAN10) to offer the edge bridge.

Figure 18-4. Edge Bridge Physical View



For detailed procedures on edge bridging configuration, see Edge Bridging: Extending Overlay Segments to VLAN. Use the same procedure for Global Manager and Local Manager.

Security in NSX Federation

You can create distributed and gateway firewall rules from the Global Manager with global, regional or local spans.

NSX Federation security provides the following benefits:

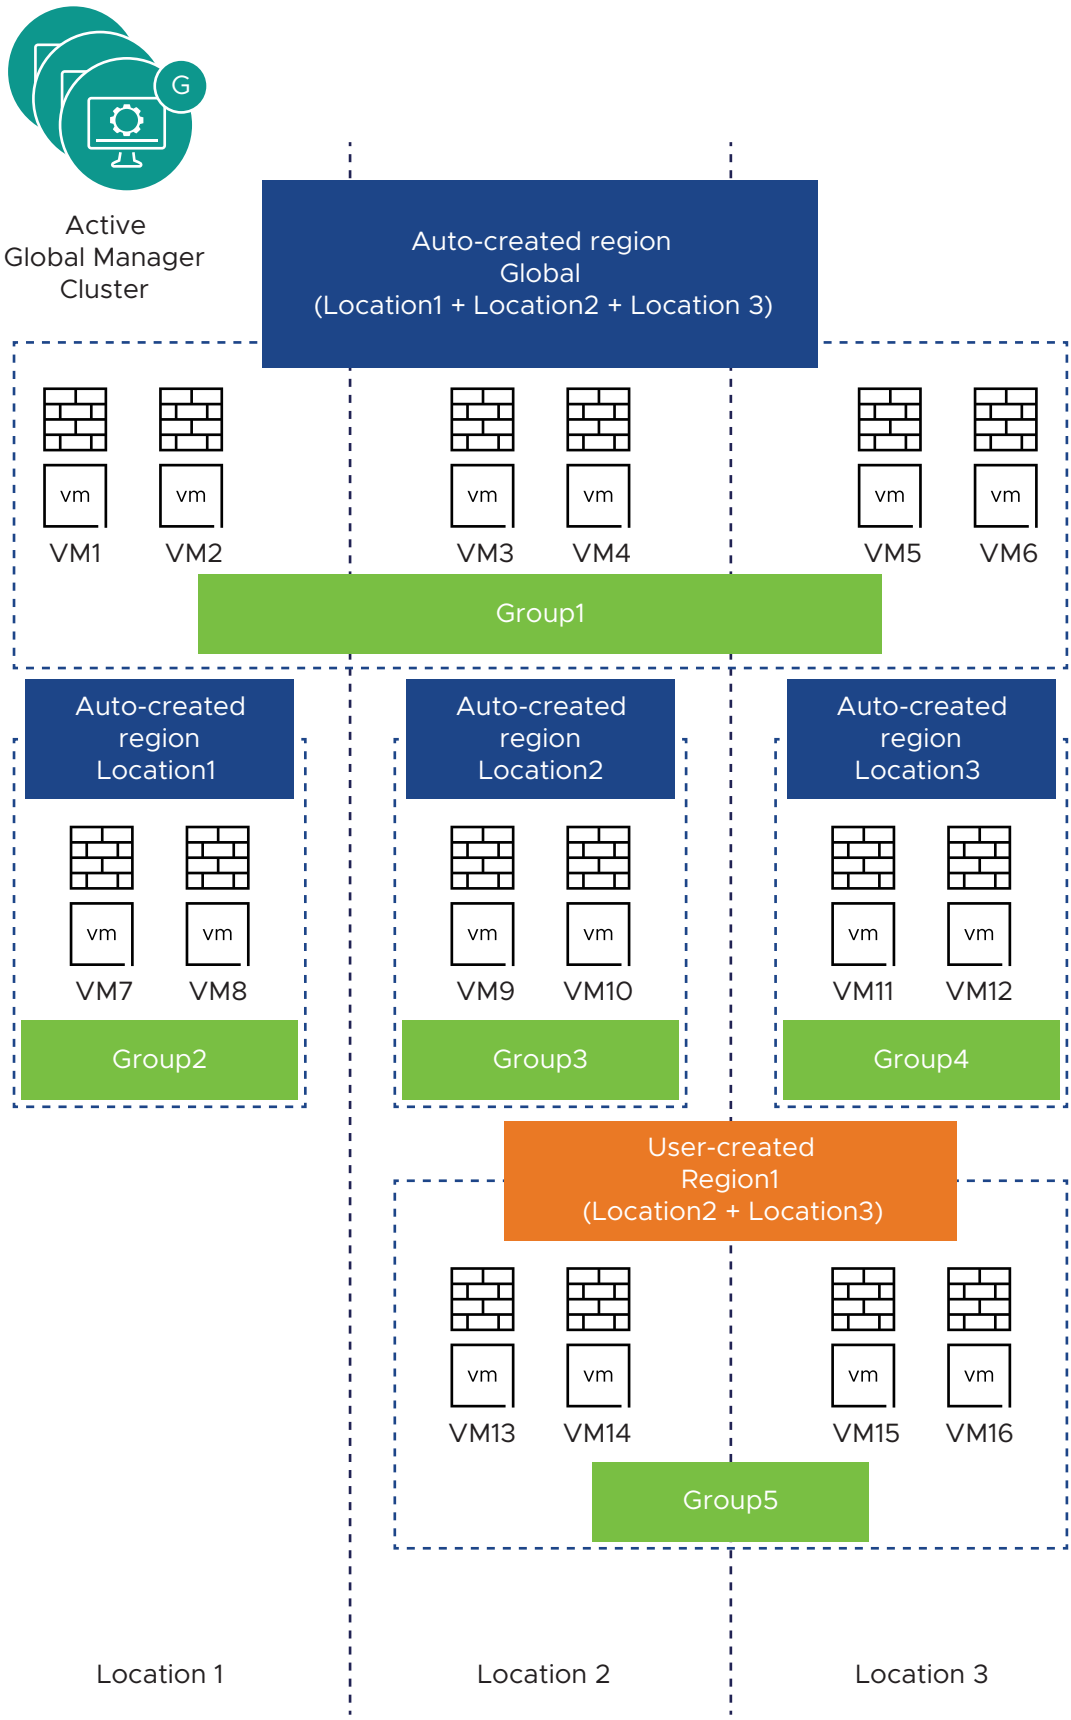
- Consistent security policy across your deployments managed using NSX Federation.
- Effective disaster recovery ensuring continuity of established security framework.
- Extension of network and security framework to another location if you are running out of compute resources in one location.

Distributed and gateway firewall policies and rules created from the Global Manager, are synced to Local Managers and appear in the Local Managers with a GM icon. You can edit rules created from the Global Manager only from the Global Manager. They cannot be edited from Local Managers.

In NSX 4.0.1.1 and later, distributed firewall is activated and deactivated with one button at the Global Manager level. Change of distributed firewall enforcement is reported at a Global Manager level, and cannot be overridden on the Local Manager level. To activate distributed firewall on the Global Manager, navigate to **Security > Distributed Firewall > Actions > General Settings** , and toggle the **Distributed Services Status** switch.

NSX Federation of Distributed Firewall (DFW) Policies and Rules

Use this example to understand the supported firewall workflows:

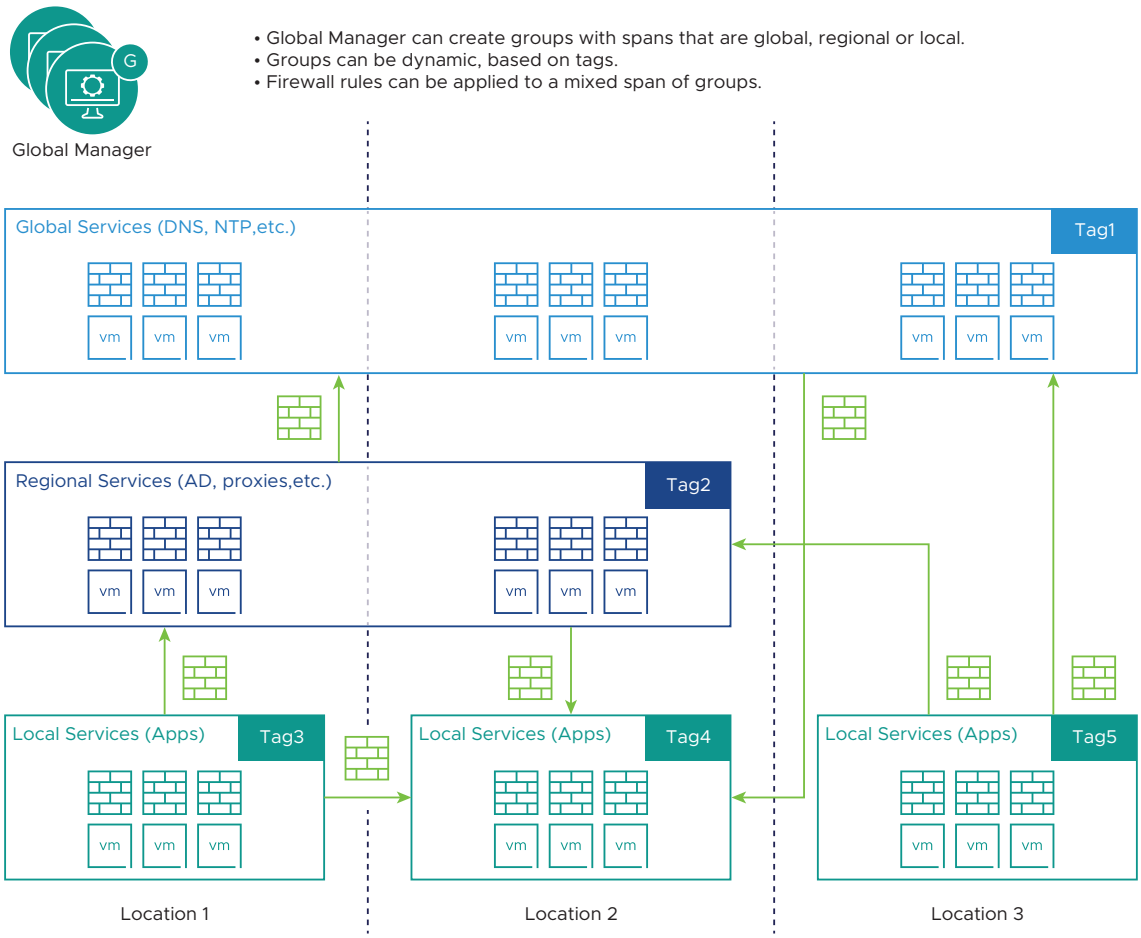

G
Active
Global Manager
Cluster
Auto-created region
Global
(Location1 + Location2 + Location 3)
vm
VM1
VM2
VM3
VM4
VM5
VM6
Group1
Auto-created
region
Location1
Auto-created
region
Location2
Auto-created
region
Location3
VM7
VM8
VM9
VM10
VM11
VM12
Group2
Group3
Group4
User-created
Region1
(Location2 + Location3)
VM13
VM14
VM15
VM16
Group5
Location 1
Location 2
Location 3

- In the example, the Global Manager has three Local Managers registered with it, named: *Location1*, *Location2* and *Location3*.
- The Global Manager auto-creates the following regions:
  - *Global*
  - *Location1*
  - *Location2*
  - *Location3*
- You create a customized region named: **`Region1`** that includes Local Managers *Location2* and *Location3*.
- You create the following groups:
  - **`Group1`**: Region *Global*.
  - **`Group2`**: Region *Location1*.
  - **`Group3`**: Region *Location2*.
  - **`Group4`**: Region *Location3*.
  - **`Group5`**: Region **`Region1`**.

DFW Policies and Rules

The following use cases are supported:

- **Group Span**: You can create groups in the Global Manager with a global, local or regional span. See Create Groups from Global Manager .
- **Dynamic Groups**: You can create groups based on dynamic criteria, such as tags.
- **DFW Policy Span**: DFW policies can be applied to a global, regional or local span.
- **DFW Rule's Source and Destination Groups**: Either all the groups in the source field or all the groups in the destination field must match the DFW policy's span. The system auto-creates groups in locations that are outside the policy's span.


• Global Manager can create groups with spans that are global, regional or local.
• Groups can be dynamic, based on tags.
• Firewall rules can be applied to a mixed span of groups.
G
Global Manager
Global Services (DNS, NTP,etc.)
Tag1
Regional Services (AD, proxies,etc.)
Tag2
Local Services (Apps)
Tag3
Local Services (Apps)
Tag4
Local Services (Apps)
Tag5
vm
Location 1
Location 2
Location 3

Refer to the table for examples of valid and invalid source and destination groups in DFW rules:

Table 18-7. Valid Source and Destination for a DFW rule based on the DFW Policy's Span

DFW Policy Span (Applied To)	Scenarios supported in DFW rules
*Global*From the example, this region contains the following groups: ■ **`Group1`**	For a DFW policy with the span of *Global* region, all groups are allowed in the DFW rule's source and destination. Following are some typical scenarios that are supported, using our example: ■ **Source**: **`Group2`**; **Destination** **`Group3`** ■ **Source**: **`Group3`**; **Destination** **`Group4`** ■ **Source**: **`Group4`**; **Destination**: *Any* ■ **Source**: **`Group1`**; **Destination** **`Group2`**.
Location1 : auto-created region for the Local Manager in location 1. From the example, this region contains the following groups: ■ **`Group2`**	For a DFW policy with the span of one location: *Location1* in this example, either the source or the destination group for the DFW rule must belong to *Location1*. The following scenarios are supported: ■ **Source**: **`Group2`**; **Destination** **`Group2`** ■ **Source**: **`Group3`**; **Destination** **`Group2`**. ■ **Source**: **`Group2`**; **Destination** **`Group4`**. ■ **Source** **`Group1`**; **Destination** **`Group2`**. The following is an example of unsupported group selections for this policy span. Both the source and the destination groups are outside the policy's span: ■ **Source** **`Group5`**; **Destination** **`Group3`**. ■ **Source** **`Group1`**; **Destination** **`Group3`**.
`Region1` : user-created region that spans *Location2* and *Location3*. From the example, this region contains the following groups: ■ **`Group5`**	For a DFW policy with the span of a user-created region: **`Region1`** in this example, either the source or the destination group for the DFW rule must contain locations that belong to **`Region1`**. The following scenarios are supported: ■ **Source**: **`Group5`**; **Destination** **`Group2`**. ■ **Source**: **`Group2`**; **Destination** **`Group5`**. ■ **Source**: **`Group2`**; **Destination** **`Group3`**. ■ **Source**: **`Group3`**; **Destination** **`Group4`**. ■ **Source**: *Any* ;**Destination**: **`Group5`** ■ **Source** **`Group4`**; **Destination** *Any* The following is an example of unsupported group selections for this policy span. Both the source and the destination groups are outside the policy's span: ■ **Source** **`Group2`**; **Destination** **`Group2`**. ■ **Source** **`Group1`**; **Destination** **`Group2`**. ■ **Source** **`Group1`**; **Destination** **`Group1`**.

- If a group contains segments, the span of the DFW policy must be greater than or equal to the span of the segment. For example, if you have a group containing a segment whose span is *Location1*, the DFW policy cannot be applied to region **`Region1`** because it only contains *Location2* and *Location3*.

NSX Federation of Gateway Firewall Policies and Rules

Gateway firewall rules can be applied to all the locations included in the gateway's span, or all interfaces of a particular location, or specific interfaces of one or more locations.

Note The span of the source and destination groups for gateway firewall rules must be the same as or a subset of the gateway's span on which you are creating the rule.

Table 18-8. Span Options for Gateway Firewall Rules

Gateway Firewall Rule's Span (Applied To)	Applies to
Apply rule to gateway	The rule applies to all interfaces attached to this gateway, in all locations that this gateway is stretched to.
Select a location and then select Apply rule to all Entities.	The rule applies only to the selected location.
Select a location and then select interfaces from that location. Repeat for other locations, selecting interfaces for each location that you want to apply the rule to.	The rule applies only to the selected interfaces.

Create a Region from Global Manager

Each location added to the Global Manager automatically becomes a region. You can also create customized regions.

Use regions to create focused groups for security and networking policies. Some regions are created automatically after you onboard locations in Global Manager. You can add more regions as necessary.

Note Each location can be a part of only one customized region.

The following regions are added by default:

- A Global region including all the locations added to the Global Manager.
- One region for each location added to the Global Manager.

For existing regions, you can view the following information:

- Name of the region.
- Locations included in the region.
- Groups the region belongs to.
- Security/Network policies the region is a part of.

Prerequisites

Refer to Security in NSX Federation for details on the implication of the span of regions and groups in creating and maintaining security policies and rules.

Procedure

1 Select **Inventory > Regions**.

2 Click **Add Region**.

3 Provide the following information:

Option	Description
Name	Provide a name for the region, for example, EMEA, or APAC.
Locations	Select the locations that you want to include in this region.

4 Click **Save**.

The region with the specified locations is created.

What to do next

Create Groups from Global Manager .

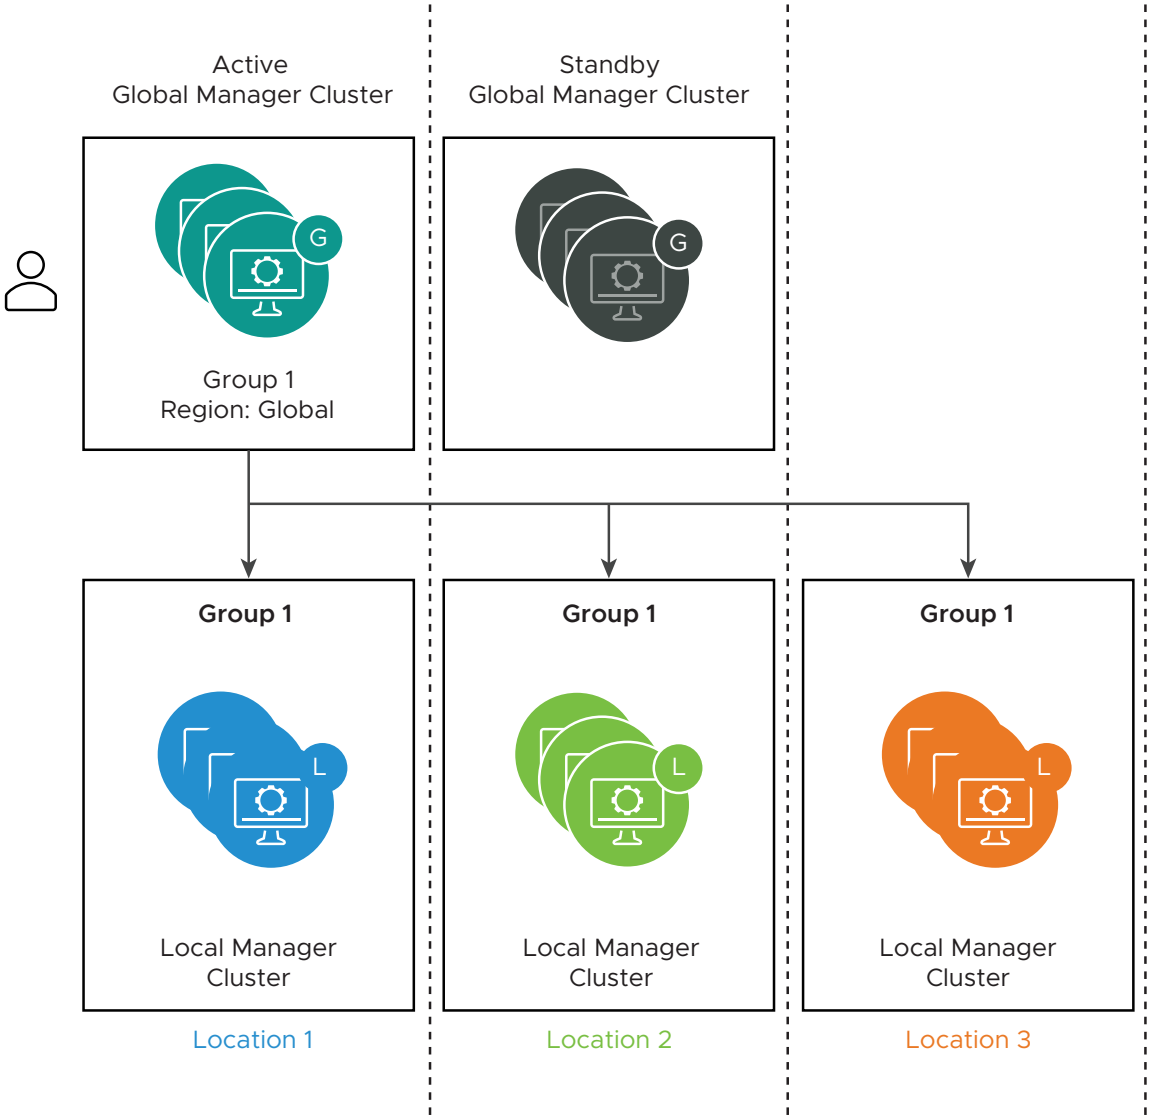
Create Groups from Global Manager

Create Groups from Global Manager that apply globally across your NSX deployments or cover selected locations or regions.

Group Span

When you create a group from the Global Manager, you select a region for the group. The group is synced with all locations in that region. A global region containing all locations, and a region for each location that has been added to the Global Manager are available automatically as regions you can select for a group's span. You can create customized regions before you create groups. See Create a Region from Global Manager.

In this example, **`Group1`** is created in the Global region, and is therefore synced with all Local Managers.



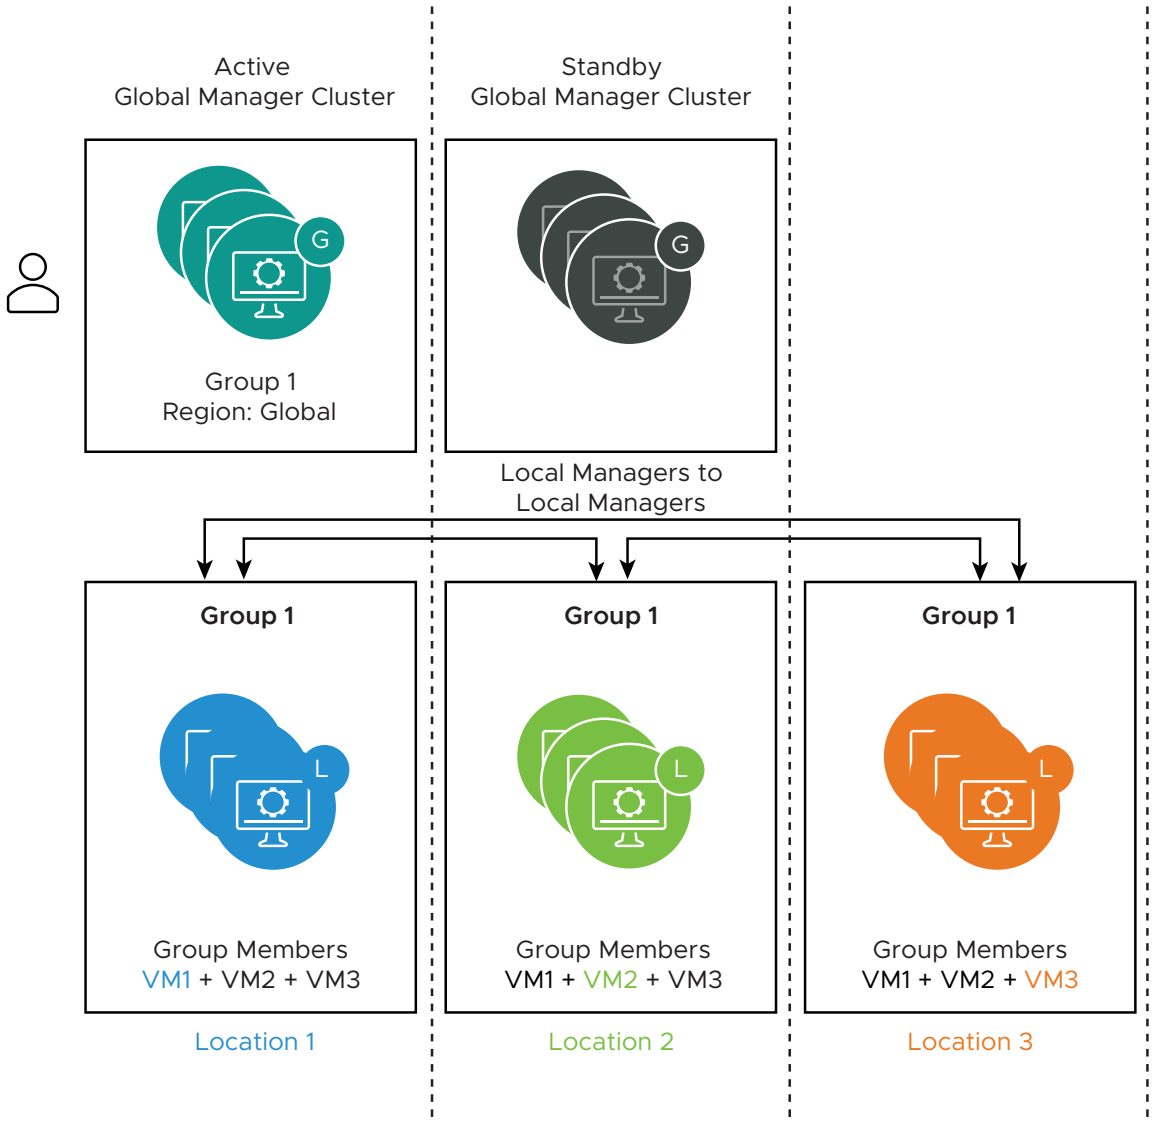
Dynamic Groups

If a group that spans more than one location has dynamic membership, you need information from each location to list the group membership.

In this example, **`Group1`** has the following members:

- **VM1** in *Location1*
- **VM2** in *Location2*
- **VM3** in *Location3*

Each Local Manager syncs its dynamic group membership with the other Local Managers. As a result, each Local Manager has a complete list of group members.



Nested Groups

For groups created from the Global Manager, you can add another group as a member if it has an equal or smaller span than the group's region.

Note If you are using NSX version 3.0.0, you can add a group as a member of another group only if the span of both the groups is exactly the same.

Extending the example using **`Region1`** that contains *Location2* and *Location3*, note the following additional configurations:

Task	Effect
From Global Manager, create **`Group-Loc2`** with region *Location2*.	■ **`Group-Loc2`** is created in Global Manager. ■ **`Group-Loc2`** is created in the Local Manager *Location2*.

From Global Manager, create group **Group-Region1** with region **Region1**. Add **Group-Loc2** as a member. This is a nested group.	■ **Group-Region1** is created in Global Manager. ■ **Group-Region1** is created in *Location2* and *Location3*. ■ **Group-Loc2** is created in Local Manager *Location3*.
From Global Manager, navigate to **Inventory > Regions** and edit **Region1** to remove *Location2*.	This action is not allowed because of the nested group **Group-Region1**.

See Add a Group for detailed steps for creating groups.

Create DFW Policies and Rules from Global Manager

You can create security policies and DFW rules to apply to multiple locations registered with the Global Manager.

Prerequisites

Ensure that you have already created any customized regions that you want to use for firewall rules. See Create a Region from Global Manager.

Procedure

1 From your browser, log in with Enterprise Admin or Security Admin privileges to a Global Manager at https://<global-manager-ip-address>.

2 Select **Security > Distributed Firewall**

3 Ensure that you are in the correct pre-defined category, and click **Add Policy**. For more about categories, see Distributed Firewall .

Note Ethernet, Emergency categories and Default Policy are not supported on Global Manager.

4 Click **Add Policy**.

5 Enter a **Name** for the new policy section.

6 Click the pencil icon next to **Applied To** to set the span of this policy.

7 In the **Set Applied To** dialog box, you can make the following selections:

- **Region**: select which Local Managers to apply the policy to. Each Local Manager is automatically added as a region. You can also create customized regions. See Create a Region from Global Manager.
- **Select Applied To**: By default, policy is applied to **DFW**, that is, the policy is applied to all the workloads on the Local Managers based on the selected region for this policy. You can also apply a policy to selected groups. Applied to defines the scope of enforcement per policy, and is used mainly for resource optimization on ESXi hosts. It helps in defining a targeted policy for specific zones, tenants and application without interfering with other policy defined for other tenants, zones & applications.

  See DFW Policies and Rules to understand how the span of the policy determines whether your DFW rule is valid or invalid.

8 To configure the following policy settings, click the gear icon:

Option	Description
TCP Strict	A TCP connection begins with a three-way handshake (SYN, SYN-ACK, ACK) and typically ends with a two-way exchange (FIN, ACK). In certain circumstances, the distributed firewall (DFW) might not see the three-way handshake for a particular flow (due to asymmetric traffic or the distributed firewall being enabled while a flow exists). By default, the distributed firewall does not enforce the need to see a three-way handshake, and picks up sessions that are already established. TCP strict can be enabled on a per section basis to turn off mid-session pick-up and enforce the requirement for a three-way handshake. When enabling TCP strict mode for a particular DFW policy, and using a default ANY-ANY Block rule, packets that do not complete the three-way handshake connection requirements and that match a TCP-based rule in this section are dropped. Strict is only applied to stateful TCP rules, and is enabled at the distributed firewall policy level. TCP strict is not enforced for packets that match a default ANY-ANY Allow which has no TCP service specified.
Stateful	A stateful firewall monitors the state of active connections and uses this information to determine which packets to allow through the firewall.
Locked	The policy can be locked to prevent multiple users from editing the same sections. When locking a section, you must include a comment. Some roles such as enterprise administrator have full access credentials, and cannot be locked out. See Role-Based Access Control.

9 Click **Publish**. Multiple policies can be added, and then published together at one time.

The new policy is shown on the screen.

10 Select a policy section and click **Add Rule**.

11 Enter a name for the rule.

12 The Source and Destination are validated based on the DFW policy's span. See DFW Policies and Rules for more information.

- If the DFW policy is applied to a location, for example, **`Loc1`**, source or destination can be either the keyword **`ANY`** or a group that belongs to **`Loc1`**.
- If DFW policy is applied to a user-created region, for example, **`Region1`** source or destination can be either the keyword **`ANY`** or a group that has the same span as **`Region1`** or spans a location in **`Region1`**.

- If DFW policy is applied to **Global**, source or destination can be anything.

Note Active Directory and IDFW are not supported for NSX Federation, that is, you cannot use these features from the Global Manager.

a In the **Sources** column, click the pencil icon and select the source of the rule.

b In the **Destinations** column, click the pencil icon and select the destination of the rule. If not defined, the destination matches any.

13 In the **Services** column, click the pencil icon and select services. The service matches any if not defined.

14 In the **Profiles** column, click the edit icon and select a context profile, or click **Add New Context Profile**. See Profiles.

15 Click **Apply** to apply the context profile to the rule.

16 By default, the **Applied to** column is set to DFW, and the rule is applied to all workloads. You can also apply the rule or policy to a selected group. **Applied to** defines the scope of enforcement per rule, and is used mainly for optimization of resources on ESXi hosts. It helps in defining a targeted policy for specific zones, tenants, and applications without interfering with other policy defined for other tenants and zones and applications.

Note You cannot select the following types of groups in **Applied to**:

- a group with IP or MAC addresses
- an Active Directory user group

17 In the **Action** column, select an action.

Option	Description
Allow	Allows all L3 or L2 traffic with the specified source, destination, and protocol to pass through the current firewall context. Packets that match the rule, and are accepted, traverse the system as if the firewall is not present.
Drop	Drops packets with the specified source, destination, and protocol. Dropping a packet is a silent action with no notification to the source or destination systems. Dropping the packet causes the connection to be retried until the retry threshold is reached.
Reject	Rejects packets with the specified source, destination, and protocol. Rejecting a packet is a more graceful way to deny a packet, as it sends a destination unreachable message to the sender. If the protocol is TCP, a TCP RST message is sent. ICMP messages with administratively prohibited code are sent for UDP, ICMP, and other IP connections. One benefit of using Reject is that the sending application is notified after only one attempt that the connection cannot be established.

18 Click the toggle button to enable or disable the rule.

19 Click the gear icon to configure the following rule options:

Option	Description
Logging	Logging is turned off by default. Logs are stored at `/var/log/dfwpktlogs.log` on ESXi hosts.
Direction	Refers to the direction of traffic from the point of view of the destination object. IN means that only traffic to the object is checked, OUT means that only traffic from the object is checked, and In/Out, means that traffic in both directions is checked.
IP Protocol	Enforce the rule based on IPv4, IPv6, or both IPv4-IPv6.
Log Label	Log Label appears in the Firewall Log when logging is enabled.

20 Click **Publish**. Multiple rules can be added and then published together at one time.

21 On each policy, click **Check Status** to view the status of rules it contains, per location. You can click **Success** or **Failed** to open the policy status window.

22 Click **Check Status** to check the realization status of policies that are applied to Transport Nodes on different locations.

Create Gateway Policies and Rules from Global Manager

You can create gateway firewall policies and rules to be applied to multiple locations or selected interfaces for particular locations, from the Global Manager.

Tier-0 or tier-1 gateways created from the Global Manager span all or a set of locations. You have a few options when applying gateway firewall rules created from the Global Manager: Gateway firewall rules can be applied to all the locations included in the gateway's span, or all interfaces of a particular location, or specific interfaces of one or more locations.

On the Local Manager rules are enforced in the following order:

1 Any rules you create from the Global Manager, that get successfully realized on the Local Manager, are enforced first.

2 Any rules that you create from the Local Manager are enforced next.

3 The last rule enforced is the default gateway firewall rule. This is the allow-all or deny-all rule applicable to all locations and all workloads. You can edit the behavior for this default rule from the Global Manager.

Procedure

1 From your browser, log in with Enterprise Admin or Security Admin privileges to the Global Manager at https://<global-manager-ip-address>.

2 Select **Security > Gateway Firewall**.

3 Ensure that you are in the correct pre-defined category. Only **Pre Rules**, **Local Gateway** and **Default** categories are supported on Global Manager. To define policy under the **Local Gateway** category, click the category name from the **All Shared Rules** tab or directly click the **Gateway Specific Rules** tab.

   Select a tier-0 or tier-1 gateway from the drop-down menu next to **Gateway**. The span of the tier-0 or tier-1 gateway you selected becomes the default span of the Gateway Firewall policy and rule. You can reduce the span but not expand it.

4 Click **Add Policy**.

5 Enter a **Name** for the new policy section.

6 (Optional) Click the gear icon to configure the following policy settings:

Settings	Description
TCP Strict	A TCP connection begins with a three-way handshake (SYN, SYN-ACK, ACK), and typically ends with a two-way exchange (FIN, ACK). In certain circumstances, the firewall may not see the three-way handshake for a particular flow (i.e. due to asymmetric traffic). By default, the firewall does not enforce the need to see a three-way handshake, and will pick up sessions that are already established. TCP strict can be enabled on a per section basis to turn off mid-session pick-up, and enforce the requirement for a three-way handshake. When enabling TCP strict mode for a particular firewall policy and using a default ANY-ANY Block rule, packets that do not complete the three-way handshake connection requirements and that match a TCP-based rule in this policy section are dropped. Strict is only applied to stateful TCP rules, and is enabled at the gateway firewall policy level. TCP strict is not enforced for packets that match a default ANY-ANY Allow which has no TCP service specified.
Stateful	A stateful firewall monitors the state of active connections, and uses this information to determine which packets to allow through the firewall.
Locked	The policy can be locked to prevent multiple users from making changes to the same sections. When locking a section, you must include a comment.

7 Click **Publish**. Multiple Policies can be added, and then published together at one time.

   The new policy is shown on the screen.

8 Select a policy section and click **Add Rule**.

9 Enter a name for the rule.

10 In the **Sources** column, click the edit icon and select the source of the rule. The source group must have the same or a subset of the gateway's span.

11 In the **Destinations** column, click the edit icon and select the destination of the rule. If not defined, the destination matches any. The destination group must have the same or a subset of the gateway's span.

12 In the **Services** column, click the pencil icon and select services. The service matches any if not defined. Click **Apply** to save.

13 In the **Profiles** column, click the edit icon and select a context profile, or click **Add New Context Profile**. See Context Profiles.

Note Context profiles are not supported for tier-0 gateways. You can apply L7 context profiles to tier-1 gateways.

14 Click the pencil icon in the **Applied to** column. In the **Applied To** dialog box:

Applied To Selection	Result
Select Apply rule to gateway	The gateway firewall rule is applied to all locations covered by the gateway's span. If you add another location to the gateway, this gateway firewall rule automatically gets applied to the location.
Select a location and then select Apply rules to all Entities	Apply this rule to all interfaces in the selected location.
Select a location and then select interfaces for that location	Apply the rule only to selected interfaces in one or more locations.

Note There is no default selection for **Applied To**. You must make a selection to be able to publish this rule.

15 In the **Action** column, select an action.

Option	Description
Allow	Allows all traffic with the specified source, destination, and protocol to pass through the current firewall context. Packets that match the rule, and are accepted, traverse the system as if the firewall is not present.
Drop	Drops packets with the specified source, destination, and protocol. Dropping a packet is a silent action with no notification to the source or destination systems. Dropping the packet causes the connection to be retried until the retry threshold is reached.
Reject	Rejects packets with the specified source, destination, and protocol. Rejecting a packet sends a destination unreachable message to the sender. If the protocol is TCP, a TCP RST message is sent. ICMP messages with administratively prohibited code are sent for UDP, ICMP, and other IP connections. The sending application is notified after one attempt that the connection cannot be established.

16 Click the status toggle button to enable or disable the rule.

17 Click the gear icon to set logging, direction, IP protocol, tag, and notes.

Option	Description
Logging	Logging can be turned off or on. You can access logs using the following NSX CLI command on NSX Edge: `get log-file syslog | find datapathd.firewallpkt` Logs can also be sent to an external syslog server.
Direction	The options are **`In`**, **`Out`**, and **`In/Out`**. The default is **`In/Out`**. This field refers to the direction of traffic from the point of view of the destination object. **`In`** means that only traffic to the object is checked, **`Out`** means that only traffic from the object is checked, and **`In/Out`** means that traffic in both directions is checked.
IP Protocol	The options are **`IPv4`**, **`IPv6`**, and **`IPv4_IPv6`**. The default is **`IPv4_IPv6`**.
Log Label	Log label that has been added to the rule.

Note Click the graph icon to view the flow statistics of the firewall rule. You can see information such as the byte, packet count, and sessions.

18 Click **Publish**. Multiple rules can be added and then published together at one time.

19 Click **Check Status** to view the realization status of policy applied to gateways through edge nodes in different locations. You can click **Success** or **Failed** to open the policy status window.

Create Drafts In Global Manager

In Global Manager, draft operations are similar to those in Local Manager.

With Global Manager, you can enforce global configuration across sites (Local Managers). One Global Manager can manage multiple Local Managers. Configurations on Global Manager are pushed to each Local Manager.

Note Local Manager can have its own specific configuration on top of global configuration.

For auto drafts and manual drafts in Global Manager, the requirements are the same as for distributed firewalls in Local Manager.

For more information on Drafts, see Firewall Drafts.

Role-Based Access Control in NSX Federation

You can configure NSX Federation role-based access control (RBAC) to restrict system access to authorized users. NSX Federation RBAC works similarly to NSX RBAC for authorized users. This topic provides some optional configuration information for RBAC on NSX Federation when it is used with specific authentication providers.

Most authentication and authorization tasks use the same procedures as described under the Authentication and Authorization section of the *NSX Administration Guide*. One exception is that the VMware Identity Manager™ (vIDM) and LDAP configuration is not synchronized from the active or the standby Global Managers (GM) to the Local Managers (LM). This requires that you configure each GM or LM (NSX cluster) separately for vIDM and LDAP. It also requires that users have the same role bindings on each NSX Federation server for seamless access.

For example, if you use NSX Federation and vIDM or LDAP authentication and want to switch between the GM and the LM server using the Location drop-down menu from the GM page, ensure you complete the following high-level tasks, so your configuration is set up properly. These configuration tasks help users that use vIDM and LDAP authentication providers avoid user permission error messages.

Task	Go To
Configure VIDM or LDAP on both the active and the standby Global Manager servers separately.	■ Integration with VMware Identity Manager/Workspace ONE Access ■ Integration with LDAP
Configure VIDM or LDAP on each Local Manager server.	■ Integration with VMware Identity Manager/Workspace ONE Access ■ Integration with LDAP
Ensure that users that want to switch between the GM and LM servers using the Location drop-down menu have the same user roles on both GM and LM servers. If the user has a role on GM, but no role on LM, users might see a permission error such as "The user does not have permission on any feature."	■ Using the Global and Local Manager Web Interfaces ■ Manage Local User Accounts ■ Create or Manage Custom Roles ■ Add a Role Assignment or Principal Identity

To ensure that the Location drop-down menu allows your user to switch between the GM and the LM servers, after you update the user roles on the LM server from read only to write or mirror the GM roles, verify that the task completes. For details, go to Using the Global and Local Manager Web Interfaces and Monitoring NSX Federation Locations.

Traceflow in Federation

Traceflow is a troubleshooting tool used to check connectivity between two NSX entities.

Traceflow in Federation is supported starting with NSX 3.2.1. The Global Manager (GM) serves as a single point to trigger traceflow all the sites. The GM gets VM inventory information from each of the sites. Only VMs (and their IP addresses that are known to the Global Manager are supported, VMs attached to ports not known to the GM are not supported. Segment ports on segments that are stretched are supported on GM. Only such ports can be sources/destination in Federation traceflow.

Procedure

1 From your browser, log in with admin privileges to the active Global Manager at https://<global-manager-ip-address>.

2 Select **Plan &Troubleshoot > Traceflow**.

3 In the Packet Information box perform the following tasks:

a Select an IPv4 or IPv6 address type.

b Select a traffic type. For IPv4 addresses the traffic type choices are: Unicast, Multicast, and Broadcast. For IPv6 address the traffic type choices are: Unicast or Multicast.

c Select a **Protocol Type**, and provide related information.

Protocol	Parameters
DHCP	Select a DHCP OP code: **Boot Request** or **Boot Reply**.
DHCPv6 (IPv6 only)	Select a DHCP message type: **Solicit**, **Advertise**, **Request**, or **Reply**.
DNS	Specify an address and select a message type: **Query** or **Response**.
ICMP	Specify an ICMP ID and a sequence.
ICMPv6 (IPv6 only)	Specify an ICMP ID and a sequence.
TCP	Specify a source port, a destination port, and TCP flags. ■ The default flag is SYN. ■ SYN cannot be combined with RST or FIN. ■ If SYN is not selected, you must select ACK or RST. ■ ACK cannot be combined with FIN, PSH, or URG.
UDP	Specify a source port and a destination port.
ARP	Select an ARP OP Code: **ARP Request** or **ARP Reply** .
NDP (IPv6 only)	For Unicast traffic, specify the **Destination IP** address. For Multicast traffic, specify the **Destination IP** and **Destination MAC** address.

4 Specify the source and destination information according to the traffic type.

Traffic Type	Source	Destination
Unicast	Select a VM or a logical port/interface. For a VM: ■ Select a VM from the drop-down list. ■ Select a virtual interface. ■ The IP address and MAC address are displayed if VMtools is installed in the VM, or if the VM is deployed using OpenStack plug-in (address bindings will be used in this case). If the VM has more than one IP address, select one from the drop-down list. ■ If the IP address and MAC address are not displayed, enter the IP address and MAC address in the text boxes. For a logical port/interface: ■ Select an attachment type: **VIF**, **DHCP**, **Edge Uplink**, or **Edge Centralized Service**. ■ Select a port.	Select a VM, a logical port/interface, or IP-MAC. For a VM: ■ Select a VM from the drop-down list. ■ Select a virtual interface. ■ The IP address and MAC address are displayed if VMtools is installed in the VM or if the VM is deployed using OpenStack plug-in (address bindings will be used in this case). If the VM has more than one IP address, select one from the drop-down list. ■ If the IP address and MAC address are not displayed, enter the IP address and MAC address in the text boxes. For a logical port/interface: ■ Select an attachment type: **VIF**, **DHCP**, **Edge Uplink**, or **Edge Centralized Service**. ■ Select a port. For IP-MAC: ■ Select the trace type (layer 2 or layer 3). For layer 2, enter an IP address and a MAC address. For layer 3, enter an IP address.
Multicast	Same as above.	Enter an IP Address. It must be a multicast address from 224.0.0.0 - 239.255.255.255.
Broadcast	Same as above.	Enter a subnet prefix length.

5 (Optional) Click **Advanced Settings** to see the advanced options.

Enter the desired values or input for the following fields, and click **Save**:

Option	Description
Hop Limit	This value should be between 2-255.
Next Header	This value should be between 2-255.
Ethertype	The default is 2048.
Timeout	The default is 30 seconds, and the range is 15-60 seconds.
Frame Size	The default is 128.
Payload Type	Select **Base64**, **Hex**, **Plaintext**, **Binary**, or **Decimal**.
Payload Data	Payload formatted based on selected type.

6 Click **Trace**.

The output includes a graphical map of the topology and a table listing the observed packets. The first packet listed has the observation type `Injected` and shows the packet that is injected at the injection point.

You can apply a filter (**All**, **Delivered**, **Dropped**) on the observations that are displayed. If there are dropped observations, the **Dropped** filter is applied by default. Otherwise, the **All** filter is applied.

The graphical map shows the backplane and router links. Note that bridging information is not displayed.

Prevent Password Lockout on Local Manager Nodes

Password lockout can occur after a Local Manager is imported to a Global Manager.

If you reset the admin password of the Local Manager, the admin account is locked out. As a result, the connectivity between the Global Manager cluster and the Local Manager cluster fails with a general system error. The Global Manager nodes in both regions must be added to an allowlist of trusted nodes for the Local Manager cluster to avoid lockdown. This can be prevented by adjusting the `lockout_immune_addresses` parameter on the Local Manager.

Procedure

1 Log in to the host that has access to your data center.

2 Update the allowlist of trusted sources on the Local Manager appliance by using the Postman PUT method.

You add the IP addresses of the Global Manager cluster to the `lockout_immune_addresses` list of theLocal Manager cluster.

a Start the Postman application in your web browser and log in.

b On the **Authorization** tab, enter the following settings and click **Update request**.

Setting	Value
Type	Basic Auth
User name	admin
Password	*nsx_t_admin_password*

c On the **Headers** tab, add a key by using the following details.

Setting	Value
Key	Content-Type
Key Value	application/json

d In the request pane at the top, send the following HTTP request.

Setting	Value
HTTP request method	GET
URL	`https://<nsx_manager_FQDN>/api/vi/cluster/api-service` (change the FQDN to your local NSX Manager FQDN)

e After a successful response ("`status: 200 OK`"), copy the returned JSON-formatted body response from the **Body** tab to a text-editor.

3 Add the `lockout_immune_addresses` information to the JSON response:

a Search for the `lockout_immune_addresses` line in the JSON response. If the line cannot be found in the JSON response, add a new line with `lockout_immune_addresses,`. Note that a "," must be added to end of the previous line.

b Add the IP addresses of all global NSX Managers (including the VIP addresses) between the brackets in the following format, leaving the quotes intact:
`"lockout_immune_addresses":[ "172.16.11.95", "172.16.11.96", "172.16.11.97", "172.16.11.98" ]`

4 Send the new security configuration to the local manager using the Postman PUT method:

a Take the previous Postman HTTP request and change the HTTP request method from GET to PUT.

b On the Body tab, paste the new JSON formatted security configuration from your code-/text editor.

c Send the new HTTP request and confirm a successful response ("status: 200 OK").

5 Repeat steps 2 through 4 for all NSX Managers.

Backup and Restore in NSX Federation

You can configure and start backups for Global Manager (GM) and each Local Manager (LM) from within the Global Manager. NSX Managers are called Local Managers if they are federated with a Global Manager.

- Log in to the active Global Manager and select **System > Backup & Restore**. Each Global Manager and Local Manager in the environment is listed. See Configure Backups for instructions.
- Backup and restore on a standby Global Manager is not supported. Backup configuration for the standby Global Manager is turned off, as this is not required. If there is an issue with a standby Global Manager, delete the standby Global Manager. Then deploy a new Global Manager and onboard it to the active Global Manager.

- You cannot restore a Local Manager from within the Global Manager. To restore a Local Manager backup, log in to the Local Manager to restore. If you create the Local Manager backup after the Local Manager is onboarded to the Global Manager, then to recover the Local Manager, you only have to restore Local Manager. After restore completes, Local Manager is automatically connected to Global Manager. See Restore a Backup for instructions.
- The system treats backup and restore operations as specific to each appliance, whether it is the Global Manager or the Local Manager you are backing up or restoring. The Global Manager's backup contains a backup of the database of that appliance only. The Local Manager contains a backup of the database and inventory of that appliance only.
- If you are restoring a Global Manager and a Local Manager, select backup timestamps of each appliance as close to each other as possible.
- After each appliance is restored, the `async replicator` service restores communication between the Global Manager and each Local Manager.

Backup Scenarios in NSX Federation

Scenario	Backup Workflow
Global Manager has any of the following changes: ■ Registering of a new LM ■ Networking configuration ■ Security configuration	Back up only the active Global Manager.

A Local Manager has any of the following changes:
■ Registering to GM
■ Networking configuration
■ Security configuration
■ Fabric changes, such as:
  ■ Host transport node added or removed (ESXi)
  ■ Edge transport nodes added or removed (VM or bare metal)
	Back up the Local Manager. You can perform this task from the Global Manager or the Local Manager.

Restore Scenarios in NSX Federation

Note, if you had a cluster of Global Manager appliances, you can only restore one node using the restore process. You must create the cluster after the restore of the first node completes.

This table uses example names for active and standby to illustrate the workflows. For detailed steps, see Chapter 26 Backing Up and Restoring NSX Manager or Global Manager. For workflows about planned and unplanned Federation NSX disaster recovery scenarios, see Disaster Recovery for Global Manager.

Scenario	Restore Workflow
Your active Global Manager, for example, GM-Loc1-Active, is lost, but you have a standby Global Manager, (GM-Loc2-Standby).	This workflow performs a switch-over to make the standby server active. 1 Go to the standby server, GM-Loc2-Standby, and make it active (resulting in GM-Loc2-Active). 2 (Optional) Recreate a standby GM-Loc1 cluster by installing the NSX Manager OVA onto a new server and name it GM-Loc1-Standby. Then make GM-Loc1-Standby active (resulting in GM-Loc1-Active). This automatically switches GM-Loc2-Active to become GM-Loc2-Standby.
Your active Global Manager, GM-Loc1-Active, is lost and you do not have a standby. For production environments, we recommend having a standby GM server.	1 Install a new GM-VM with the same IP address as the lost GM-Loc1-Active server. Use the NSX Manager OVA (resulting in GM-Loc1). Ensure you delete the active server before proceeding with this install. For instructions, see the *NSX Installation Guide*. 2 After installation, access the UI, select **System > Backup & Restore** then: a Enter the SFTP server information where the backup is stored. b Select the latest backup file. c Click on the Restore.
Your standby Global Manager, GM-Loc2-Standby, is lost.	1 On your active GM, GM-Loc1-Active, remove the current GM standby, GM-Loc2-Standby, using **System > Location Manager** . 2 Install a new GM-VM using the NSX Manager OVA (resulting in GM-Loc2). You can optionally use the same IP address as the old GM-Loc2-Standby. 3 On the active GM, GM-Loc1-Active, add the newly installed GM-VM, GM-Loc2, as standby. To add the standby after installation completes, on the active GM, select **System > Location Manager** (resulting in GM-Loc2-Standby).
A Local Manager, LM-Loc1, is lost.	Restore the Local Manager, LM-Loc1, from the Local Manager. See Chapter 26 Backing Up and Restoring NSX Manager or Global Manager. When restored, configurations from the Global Manager are synchronized with the Local Manager.
Both the Global Manager and the Local Manager are lost.	If you are restoring both the Global Manager, GM-Loc1-Active, and the Local Manager, LM-Loc1, use the latest backups of each appliance. When the Global Manager and the Local Manager are restored, the Global Manager, GM-Loc1-Active, pushes the configurations to the Local Manager, LM-Loc1. You must manually resolve any discrepancies in inventory and fabric related changes between the Local Manager and the Global Manager.

Disaster Recovery for Global Manager

In an NSX Federation environment, if you lose your active Global Manager, you can switch to the standby Global Manager.

These workflows describe the following scenario where GM denotes the Global Manager appliance.

You have a GM service to federate the network and security services across your locations. This includes:

- One GM cluster in location `Loc1` set **Active** (**`GM-Loc1`**).
- One GM cluster in location `Loc2` set **Standby** (**`GM-Loc2`**).

Planned Switchover to Standby Location

The active GM service is in Loc1 and the stretched tier-0/tier-1 primary routing services are also in Loc1.

1 Switch your GM service in `Loc2`.

   a Log in to the standby GM – **`GM-Loc2`**.

   b Select the **Active** drop-down menu and click **Make Active**.

   The system starts the process of making GM-Loc2 active. After the process completes, GM-Loc2 changes status to Active and GM-Loc1 changes status to Standby.

   If the planned failover is unsuccessful use the Force Option.

   Note If the synchronization between GM-Loc1 and GM-Loc2 is not successful for some time, you might have possible configuration data loss when GM-Loc2 is forced to be GM Active.

   c To check if GM-Loc1 and GM-Loc2 are synchronizing, run `/etc/init.d/corfu-log-replication-server status` in an SSH root shell command on both GM-Loc1 and GM-Loc2. If synchronization is not running, contact technical support.

2 (Optional) Switch primary stretched tier-0/tier-1 `Loc1` to `Loc2`.

   a To move the stretched tier-0 and tier-1 gateways to the secondary site, follow the network recovery workflow. See Network Recovery for Local Managers.

   b Recover the compute VMs using your preferred method. For example, use VMware Site Recovery Manager from `Loc1` to `Loc2`.

Unplanned Switchover to the Standby Location

If the site `Loc1` becomes unavailable, **`GM-Loc1`**, **`LM-Loc1`**, as well as `Loc1` Edge Nodes stop responding. Proceed with a forced failover.

1 Recover your GM Service in `Loc2` .

   a Log in to the standby GM – **`GM-Loc2`**.

   b Select the **Actions** drop-down menu and click **Make Active**.

   The system starts the process of making GM-Loc2 active. After the process completes, GM-Loc2 changes status to **Active** and all LMs (other than LM-Loc1 unavailable) synchronizes with GM-Loc2 from the last synchronized configuration received from GM-Loc1.

If **GM-Loc1** is online after **GM-Loc2** changes to active, the status of **GM-Loc1** is NONE. To make **GM-Loc1** standby:

1. Log in to the active GM – **GM-Loc2**.
2. From the tile where GM-Loc1 shows the status of NONE, select the **Actions** drop-down menu and click **Make Standby**.

2. Move your primary stretched tier-0/tier-1 **Loc1** to **Loc2**.
   a. Follow the network recovery workflow to move stretched tier-0 and tier-1 gateways to the secondary site. See instructions at Network Recovery for Local Managers.
   b. Recover the compute VMs using your preferred method. For example, use VMware Site Recovery Manager from to `Loc2`.

See section 4.4 titled *Disaster Recovery* in the NSX-T Data Center Multi-location Design Guide for more details. For details about using Site Recovery Manager, see Working with VMware Site Recovery Manager and Multisite Environments.

Working with Site Recovery Manager and Federation

You can use Site Recovery Manager (SRM) with NSX Federation for disaster recovery use cases.

SRM supports the following workflows with NSX Federation:

- NSX Federation Global Manager (GM) VMs support full and test recovery of GM VMs (supported with or without NSX Federation management cluster VIP).
- Compute VMs support full and test recovery of compute VMs. Recovered VMs in the disaster recovery site have their NSX tags and firewall rules based on these NSX tags or not such as IP addresses and VM names.

To ensure that groups and firewall rules replicate at the disaster recovery location during recovery, the NSX Local Manager managing the disaster recovery location must have the NSX tags present at recovery time.

SRM does not replicate the NSX VM tags to the Local Managers managing disaster recovery location.

Prior to NSX Federation 3.2, had no support for replication of VM tags between Local Managers. As a result, NSX did not replicate any security based on VM tags to recovered VMs. Security not based on VM tags, for instance, IP addresses or VM names, gets applied to recovered VMs.

To configure VM tag replication across Local Managers run the following Global Manager API:

```
PUT https://{{gm}}/global-manager/api/v1/global-infra/vm-tag-replication-policies/policy1
{
    "display_name":"vm tag replication policy Paris to London",
    "description":"vm tag replication policy1",
    "protected_site": "/global-infra/sites/LM_Paris",
    "recovery_sites": [
        "/global-infra/sites/LM_London"
    ],
```

```
    "groups":[
        "/global-infra/domains/default/groups/Web-VM-Group",
        "/global-infra/domains/default/groups/DB-VM-Group"
    ],
    "vm_match_criteria": "MATCH_BIOS_UUID_NAME"
```

In the previous Global Manager configuration, LM_Paris sends the tag information of the VMs for the BIOS UUID of the VMs in the groups Web-VM-Group + DB-VM-Group to LM_London. Before the recovery of the London VMs by SRM, LM_London does not have the VMs with the BIOS UUID and the VMs are not visible in LM_London yet. However, when SRM recovers the VMs in London, LM_London sees those VMs with the BIOS UUID and applies their NSX tags on them. The VMs get their security based on NSX tags.

Note `vm_match_criteria` has two possible values MATCH_BIOS_UUID_NAME or MATCH_NSX_ATTACHMENT_ID. At the recovery, SRM copies both so any configuration is valid with SRM. However, if another product completes VM replication and copies one, but not the other value, then configure GM with the appropriate vm_match_criteria value.

To get details on VM tag replication across Local Managers run the following Global Manager API :

```
GET https://{{gm}}/global-manager/api/v1/global-infra/vm-tag-replication-policies
```

The output returns something similar to:

```
{
  "protected_site": "/global-infra/sites/LM_Paris",
  "recovery_sites": [
    "/global-infra/sites/LM_London"
  ],
  "vm_match_criteria": "MATCH_BIOS_UUID_NAME",
  "groups": [
    "/global-infra/domains/default/groups/Web-VM-Group",
    "/global-infra/domains/default/groups/DB-VM-Group"
  ],
  "resource_type": "VMTagReplicationPolicy",
  "id": "policy1",
  "display_name": "vm tag replication policy Paris to London",
  "description": "vm tag replication policy1",
  "path": "/global-infra/vm-tag-replication-policies/policy1",
  "relative_path": "policy1",
  "parent_path": "/global-infra",
  "unique_id": "9ee18586-5480-41d9-8223-690c9226d763",
  "marked_for_delete": false,
  "overridden": false,
  "_create_time": 1638413861377,
  "_create_user": "admin",
  "_last_modified_time": 1638413861377,
  "_last_modified_user": "admin",
```

```
  "_system_owned": false,
  "_protection": "NOT_PROTECTED",
  "_revision": 0
}
```

NSX supports only one entry from recovery sites. For details, see the vm-tag-replication-policies/policy-name API in the *NSX Global Manager REST API Guide*.

Network Recovery for Local Managers

If a Local Manager is lost, you can recover networking configurations from it using the auto-detected Network Recovery option in the Global Manager.

You must have at least one stretched tier-0 or tier-1 gateway set up designating a Location Manager as primary. The loss of this primary Location Manager for the tier-0 or tier-1 gateway triggers the option of network recovery in the Global Manager.

- The Global Manager detects the loss of connection and prompts you to perform **Network Recovery**.
- In the first step of recovery, you recover the tier-0 gateway. You can change the preferred primary location if you want it to be different from the one you set in the fallback preference.
- In the second step, you select a preferred primary location for tier-1 gateways that have a subset of the span of the locations covered by the tier-0 network. The preferred primary location for such tier-1 gateways would be different from tier-0 gateways and you must either accept the fallback preference established by the tier-0 gateway, or elect not to move the gateway.
- In the final step, you can view the list of networking constructs that cannot be recovered because they do not have a secondary location configured.

Note If you have a tier-0 and tier-1 gateway set up using a Location Manager as primary, but the tier-0 and tier-1 gateway do not have any services attached to them, for example, tier-0 and tier-1 without NAT and firewall, then the data plane traffic still works after the loss of the primary Location Manager. For tier-0/tier-1 configuration without service, Network Recovery is not mandatory for the recovery of data plane, even though the Network Recovery option appears in the Global Manager.

Procedure

1. From your browser, log in with admin privileges to the active Global Manager at https://<global-manager-ip-address>.
2. Select **System > Location Manager**.
3. A banner appears on this page noting the location that is down. Click **Network Recovery** on the banner and start the workflow for **Location Disaster Recovery** in the following steps.

4 **Tier-0 Gateways**: For each tier-0 gateway that has the failed location set as primary, you have the option to select a new primary location. This new primary location can be different from the fallback preference you elected when creating the tier-0 gateway. You can also elect to not move the tier-0 gateway. Click **Apply Configuration** for each tier-0 gateway after selecting a new primary location or retaining the priority set earlier.

5 Click **Next**.

6 Tier-1 gateways are listed for recovery only if their span differs from the span of the tier-0 gateway. If tier-1 gateways follow the same span as the tier-0 gateway, the same locations are selected to be primary as for tier-0 gateways. For a different span, you can either select a different location as primary or elect to not move the tier-1 gateway at all.

7 After you make your selections for each tier-1 gateway, click **Accept** and **Next** to proceed.

8 Under **Single Location Entities** you can see a list of tier-0 and tier-1 gateways that cannot be moved to a new primary location because they exist only in the failed location. Click **Next** to proceed.

Results

The stretched tier-0 and tier-1 gateways are moved to the new location which that you designated as primary.

See section 4.4.2 titled *Data Plane Recovery* in the NSX-T Data Center Multi-location Design Guide for more details.

Multi-tenancy

19

NSX supports multi-tenancy where you can configure multiple tenants on a single NSX deployment. Multi-tenancy enables you to isolate security and networking configuration across tenants.

NSX users across tenants can access the same NSX Manager to provision networking and security objects on shared hosts.

This chapter includes the following topics:

- Projects

Projects

A project in NSX is analogous to a tenant. By creating projects, you can isolate security and networking objects across tenants in a single NSX deployment.

Let us assume that an organization named "ACME Enterprise" has NSX deployed at its site. This organization currently has all its infrastructure, networking, and security configurations in the default space, which is owned by the enterprise administrator. You will learn more about the default space later in this documentation.

This organization has the following objectives:

- Isolate networking and security configurations for its three departments: Sales, Marketing, and Operations.
- Delegate the tasks of creating and managing networking and security configurations for each department to a specific set of NSX users, and avoid giving these users visibility to all the objects in the system.
- By default, allow workload VMs within a department to communicate only with other workload VMs (including the DHCP server) in the same department.
- By default, block communication with workloads outside the department. If such a communication is required, the system must allow adding new rules or modify existing rules in the default security policy.

To achieve these objectives, ACME Enterprise can implement multi-tenancy in their NSX deployment by using the Project feature. For example, it can create three projects with these names:

- ACME-Sales
- ACME-Marketing
- ACME-Operations

In a multi-tenancy deployment, users in each project have access to objects that they create in their project and can consume objects (in a read-only mode) that the enterprise administrator has shared with their project from the default space. You will learn more about Resource Sharing feature later in this documentation.

Note Setting up multi-tenancy in your NSX deployment is optional and its implementation has no impact on your existing NSX configuration.

Multi-tenancy Policy Data Model

The NSX Policy data model is hierarchical and has two system-created branches:

- `/infra` branch is managed by the enterprise administrator. The objects under this branch are displayed in the **Default** view on the UI.

  User roles other than the enterprise administrator also exist in the `/infra` branch. The users in this branch are not tied to any specific projects. This documentation refers to such users as "system-wide" users. These users can configure a subset of objects under the `/infra` branch.

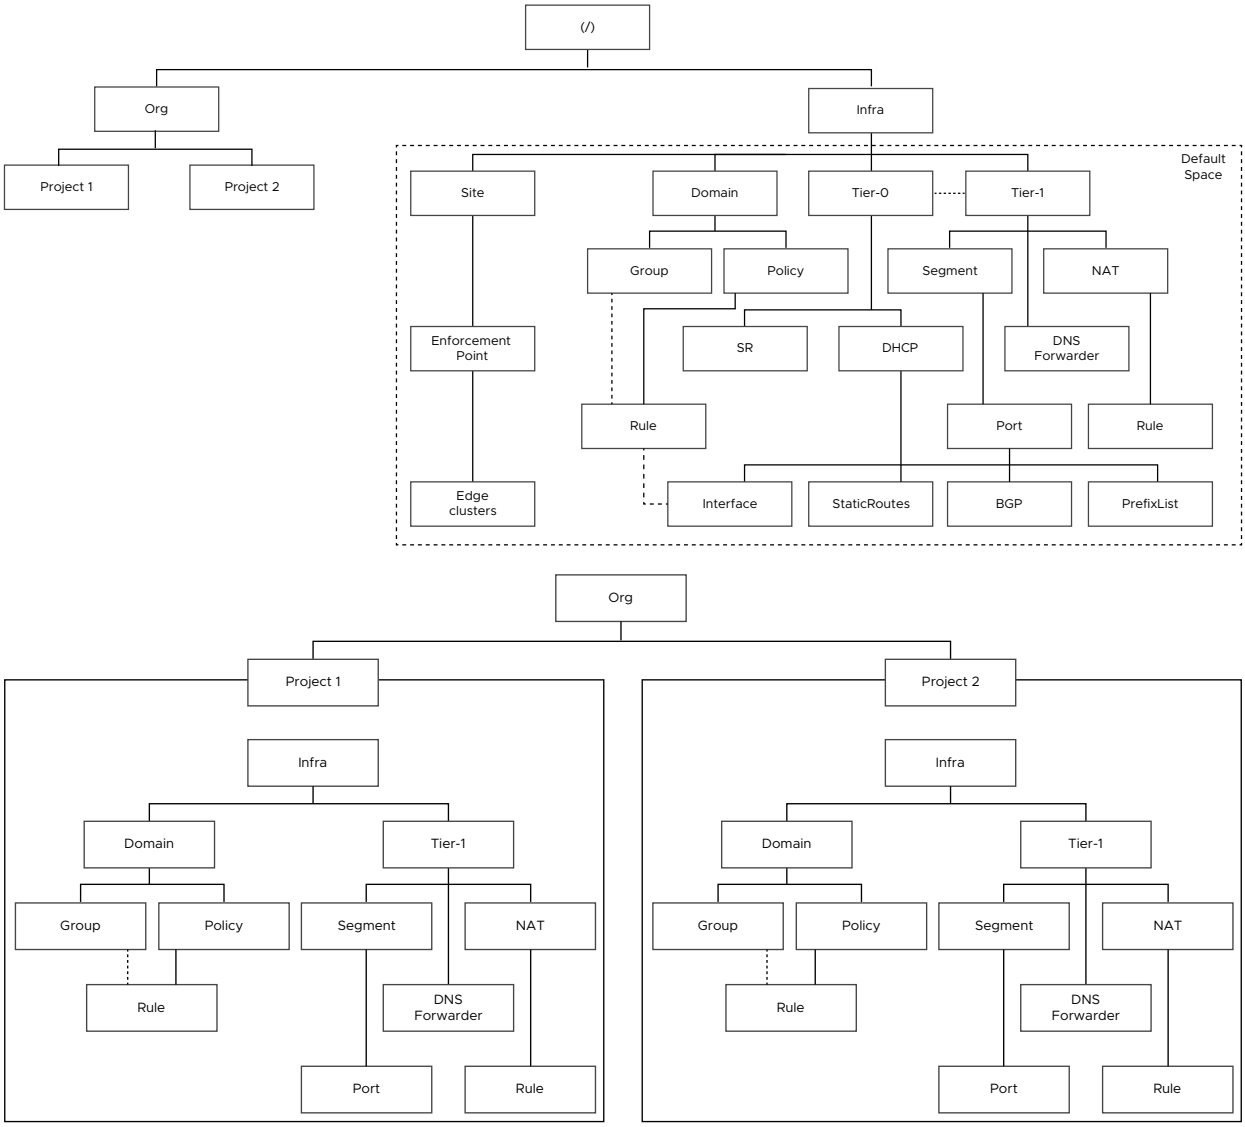
  System-wide users have access to all the objects in the system. That is, they have access to objects inside the **Default** view (`/infra` branch) and inside the projects.

  At times, this documentation uses the term "default space" to refer to objects under the **Default** view. In other words, the terms "default space" and "**Default** view" are used interchangeably. They both mean the same. To learn more about the **Default** view, see the *Overview of the Default View (Default Space)* subsection later in this documentation.

- `/orgs/default` branch holds the multi-tenancy objects. Every project has its own space to host objects that it owns.

The following diagrams illustrate the data model for multi-tenancy. The first figure shows the default space and two projects under the org. The next figure shows the hierarchy of objects in both the projects. Under the org, projects 1 and 2 have their own hierarchy of NSX networking and security objects that are created inside the project. Objects created inside a project are owned by that project.

Tier-0 gateways and edge clusters are owned by the default space, and they can be allocated to projects under the org. You cannot create tier-0 gateways and edge clusters inside a project.



Default Org

An NSX deployment has one default org. You cannot create, modify, or delete the default org. The org object is created by the system at startup. The tier-0 gateways and edge clusters in the system can be allocated to projects under the org.

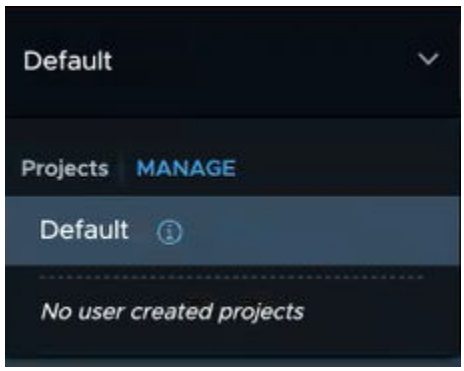
The org object is created by the system with the following identifier:

`/orgs/default`

The org object is not visible in the UI.

Understanding the Project Switcher Menu

The project switcher is a drop-down menu on the application bar, which is at the top of the NSX Manager UI, as shown in the following screen capture.



Overview of the Default View (Default Space)

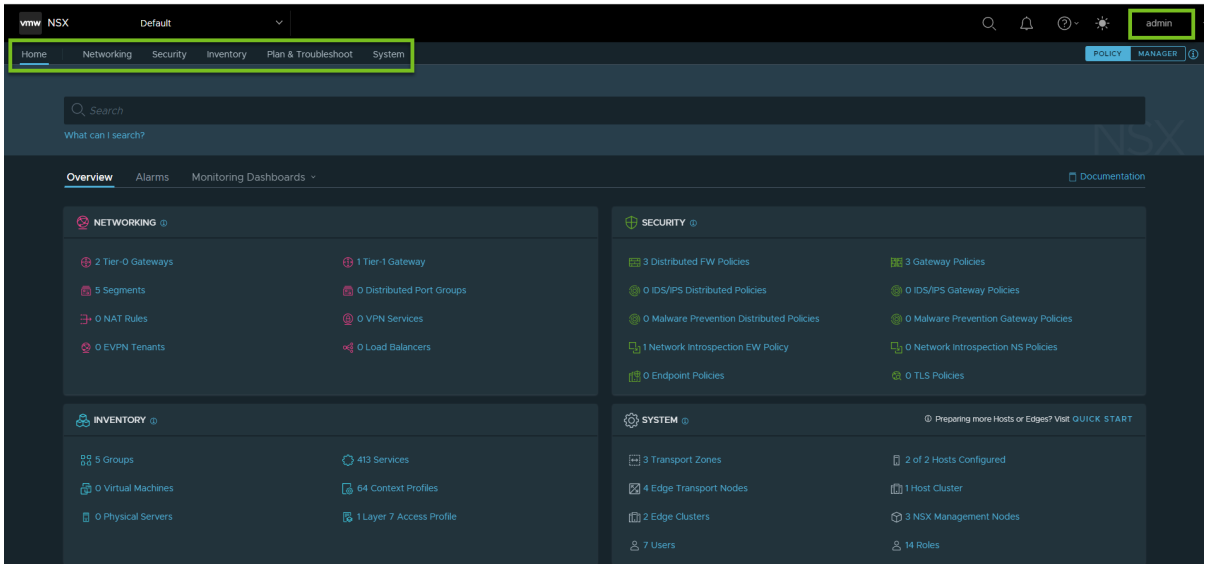
When you log in to NSX Manager for the first time, the project switcher shows only the **Default** view. No user-created projects exists in the system, as shown in the preceding screen capture.

The **Default** view is visible to the enterprise administrator and to the other system-wide user roles that are not assigned to any specific projects. This view contains:

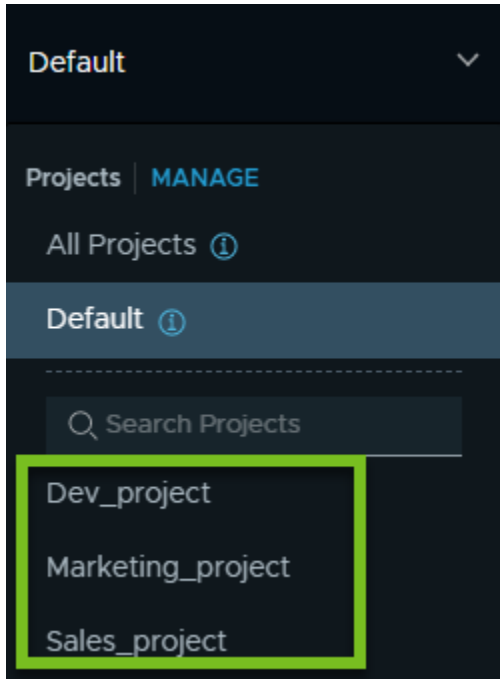
- All the objects in your NSX fabric, such as hosts, tier-0 gateways, edge clusters, transport zones, and so on.
- Global user management objects.
- Objects under the `/infra` space of the hierarchical Policy data model, such as tier-1 gateways, segments, groups, firewall policies, and so on.
- Resource shares

In short, the **Default** view contains NSX objects that do not belong to any project.

For example, when an enterprise administrator logs into NSX Manager, the **Default** view is displayed, as shown in the following screen capture. Observe that all the tabs are displayed in the UI. The **Overview** page shows a high-level summary of the objects in the default space.

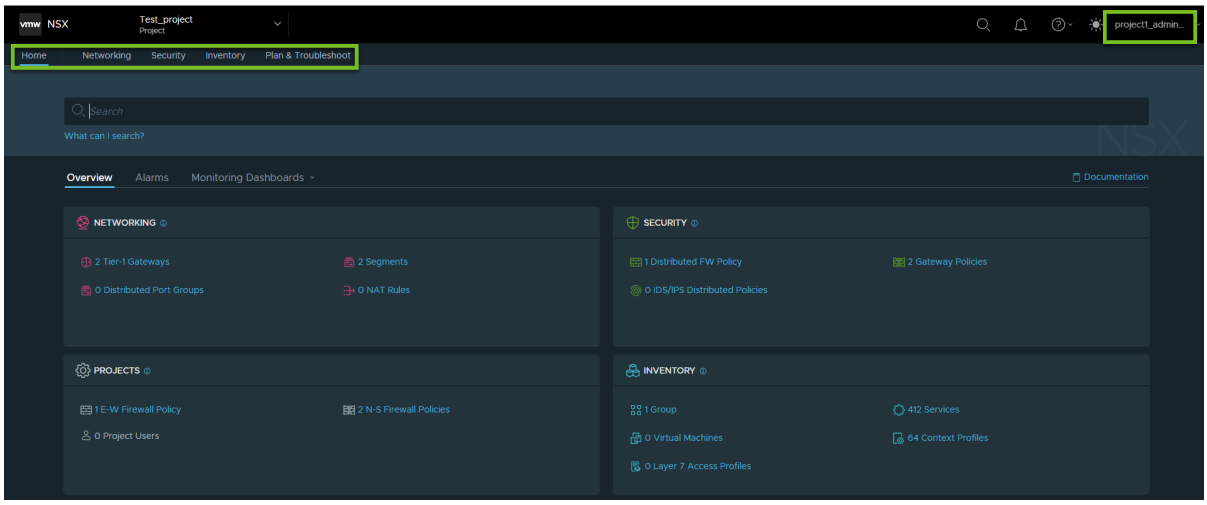


After one or multiple projects are created in your NSX deployment, these projects are displayed in the project switcher drop-down menu, as shown in the following screen capture.



An enterprise administrator can view all the projects in the system. Other system-wide user roles, such as an auditor, can also view all the projects in the system. Users that are assigned to specific projects with roles, such as project admin, security admin, network admin, security operator, and network operator can view the projects that they have access to.

For example, when a project administrator logs into NSX Manager, the project-specific view is displayed, as shown in the following screen capture. Observe that the **System** tab is not visible to the project administrator. The **Overview** page displays a high-level summary of the objects that are created in the project.



When you upgrade from an earlier version of NSX to NSX 4.1, the default space will host all the existing infrastructure, networking and security configurations. You can continue to use the default space for all the networking and security requirements in your organization. The properties of the existing networking and security objects or the path of those objects are not modified. Creating projects is optional.

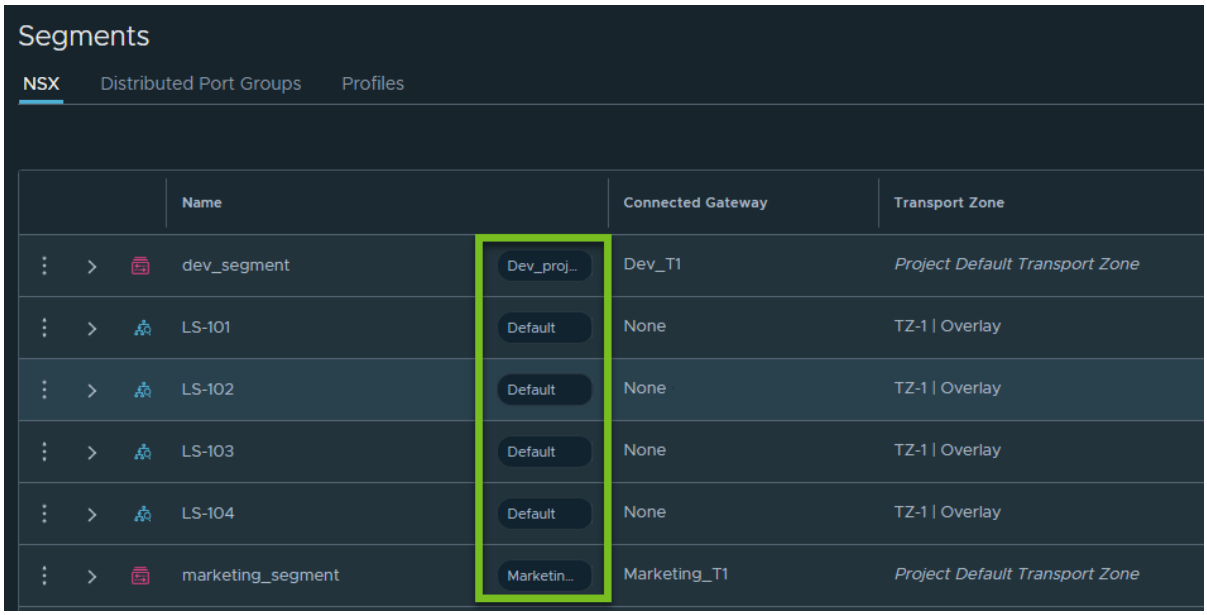
If you created projects in your NSX 4.0.1.1 deployment by using the NSX API, and then upgraded to 4.1, the **Default** view and project views are listed in the project switcher drop-down menu. You can switch between project views and the **Default** view to view the objects under each of them. In addition, the project switcher drop-down menu also shows the **All Projects** view, which is described in the following section.

Overview of the All Projects View

- The **All Projects** view is available to all the system-wide user roles that are not assigned to any specific project. That is, this view is available to all the users that have access to the default space. For example, enterprise administrator, auditor, and so on.
- This view is available in the project switcher drop-down menu only after at least one project is added in your NSX deployment.
- This view displays networking and security configurations across all projects, including the default space.
- The objects in this view are displayed in a **read-only** mode.

In the **All Projects** view, the networking and security objects display a pill-shaped icon next to the object name to indicate whether the object is owned by the default space or a project.

For example, the following screen capture shows the list of segments on the **Segments** page. The pill-shaped icons that are highlighted in the green box indicate who owns each segment – default space or the project.



Overview of a Project

Under the org, you can create a project for each tenant that you want to manage from your NSX environment. Projects are created under `/orgs/default` to support independent sets of networking and security configurations for each tenant.

Projects offer the Policy hierarchical model in a specific isolated context. Project configurations are set up under `/orgs/default/projects/<project-id>/infra`

When a project is added, some default objects are created automatically in the project, such as default domain, default groups, and default security policies.

You cannot create tier-0 gateways and edge clusters inside a project. The tier-0 gateways and edge clusters from the default space can be allocated to projects under the org. To learn about the default space, see the *Overview of the Default View (Default Space)* subsection in this documentation.

Each project can optionally have its own tier-1 gateways, which must be configured in the project. In other words, the tier-1 gateways must be owned by the project. A project cannot use the tier-1 gateways that are configured in the default space.

Project Creation

You must have the **Enterprise Admin** role to add projects in your system.

To learn more about adding projects, see Add a Project.

User Role Assignment

The enterprise administrator is the super user in your system who can create projects in your NSX deployment, and this user has access to all the objects within and outside projects.

A project administrator has full access to all the objects within a project.

Users with different roles can be assigned to a project. Users can configure networking and security objects inside the project, as determined by permissions that are configured for each role. To learn more about user role assignment, see Add User Role Assignments for a Project.

Resource Sharing

When a project is added in the org, no user-created objects exist in the project. A new project has access to the system-defined NSX objects that are available by default to all the projects in the org. These system-defined objects are included in the default share, which is created automatically when NSX is deployed. Examples of system-defined NSX objects are Services, BFD Profiles, App IDs, Context Profiles, and many more.

A subset of objects can be shared from the default space to specific projects or to all the projects in the org. This object sharing is done by creating resource shares. Each resource share is identified by a unique name. Project users can consume the shared objects in their projects to configure groups, firewall rules, and so on, to meet their networking and security requirements. To learn more about the default share and sharing resources with projects, see:

- Resource Sharing
- Share Resources from the Default View

Quota Assignment

After adding a project in the org, one or multiple quotas can be added to define a maximum limit for the number of objects of a specific type that users can create in that project. If a quota is not applied to a project, system-defined maximum limits apply.

The system-defined limits apply to the number of objects of a specific type across all the projects, including the objects in the default space. The number of objects of a specific type across the entire system must not exceed the maximum limits that are mentioned on the VMware Configuration Maximums portal.

To learn more about quotas, see Add Quota to a Project.

NSX Features Available Under a Project

A subset of NSX features are available for consumption under projects. To learn more, see Features Available for Consumption Under Projects.

Hierarchical API for Projects

Hierarchical API provides users a way to create an entire or a part of an intent by invoking a single API call. The input to the API is expressed in a tree format. Each node in tree can have multiple children of different types.

Hierarchical API is supported for projects. All project users can invoke the following Patch API at `org-root` to create, update, or delete an entire intent hierarchy, or a part of an intent hierarchy.

```
PATCH https://<nsx-mgr>/policy/api/v1/org-root
```

Example Request Body:

```
PATCH https://<nsx-mgr>/policy/api/v1/org-root
{
    "resource_type": "OrgRoot",
    "children": [{
        "resource_type": "ChildResourceReference",
        "id": "default",
        "target_type": "Org",
        "children": [{
            "resource_type": "ChildProject",
            "Project": {
```

```
                "id": "project-1",
                "resource_type": "Project",
                "display_name": "Test HAPI Project",
                "site_infos": [{
                    "edge_cluster_paths": [
                        "/infra/sites/default/enforcement-points/default/edge-clusters/ec1"
                    ],
                    "site_path": "/infra/sites/default"
                }],
                "tier_0s": [
                    "/infra/tier-0s/vmware"
                ]
            }
        }]
    }]
}
```

Role based access control is applied to the response payload of this Patch API.

For a detailed information about the `org-root` schema, see the *NSX API Guide*.

Only an enterprise administrator can invoke the Get API call to read the `org-root` properties.

```
GET https://<nsx-mgr>/policy/api/v1/org-root
```

Add a Project

Projects help you to isolate networking and security configurations across tenants in a single NSX deployment.

Prerequisites

You must be assigned the **Enterprise Admin** role.

Procedure

1. From your browser, log in to an NSX Manager at `https://`*`nsx-manager-ip-address`*.
2. Expand the project switcher drop-down menu from the application bar at the top, and then click **Manage**.
3. Click **Add Project**.

4 Configure the project settings.

a Enter a name for the project.

b (Optional) From the **Tier-0/VRF Gateways** drop-down menu, select a tier-0 or a tier-0 VRF gateway that the workloads in this project can use for north-south connectivity with the physical network outside NSX.

You can select multiple gateways, if required. If no gateway is selected, the workloads in the project will not have north-south connectivity.

A tier-0 or a tier-0 VRF gateway can be assigned to multiple projects. That is, allocating them to one project (say project 1) does not prevent you from allocating them to other projects (say project 2 and project 3).

c (Optional) From the **Edge Clusters** drop-down menu, select an edge cluster to associate with this project.

The edge clusters that you select does not imply that the tier-0 gateways will run on them. The selected edge clusters can be consumed inside the project in the future. For example, the edge clusters can be consumed for running centralized services such as NAT, Gateway Firewall, Gateway DHCP, and so on, that you configure on the tier-1 gateways inside the project.

If you plan to use tier-1 gateways in the project only for distributed routing, but not for centralized services, you might want to skip this step.

An edge cluster can be assigned to multiple projects. That is, allocating an edge cluster to one project (say project 1) does not prevent you from allocating it to other projects (say project 2 and project 3).

Note The edge cluster must belong to the default overlay transport zone.

d (Optional) In the **Short log identifier** text box, enter a string that the system can use to identify the logs that are generated in the context of this project.

Currently, the short log identifier is applied only to the security log files.

This identifier must be unique across all the projects in your NSX deployment. The identifier must not exceed eight alphanumeric characters. If it is not specified, the system autogenerates it when you save the project. After the identifier is set, you cannot modify it.

e (Optional) Enter a description for the project.

5 Click **Save**.

Add User Role Assignments for a Project

After adding a project in your org, you can assign roles to users in the project. These users can then start configuring networking or security objects in the project.

In a project, you can assign the following roles to users:

- Project Admin
- Security Admin
- Network Admin
- Security Operator
- Network Operator

A project administrator has full access to all the networking and security objects inside a project. The other project-specific user roles have limited access to objects inside the project as determined by the permissions of those roles.

Note By default, only an enterprise administrator can add user role assignments in projects. A project administrator does not have permissions to add user role assignments in projects, unless an enterprise administrator configures the **Project Admin** role to do user role assignments.

If the **Project Admin** role is allowed to do user role assignments, it might be perceived as introducing security risks in some NSX environments because it allows a project administrator to configure role assignments for any user in the system. Therefore, the default behavior is to allow only enterprise administrators to add user role assignments in projects.

An enterprise administrator can do the following steps to grant permissions to the **Project Admin** role for adding user role assignments:

1 In the **Default** view, navigate to **System > User Management > Roles**.

2 Next to the **Project Admin** role, click the  icon , and then click **Allow Role Assignments**.

For user authentication and authorization, NSX multi-tenancy supports the following identity sources:

- Local users (for example, guestuser1, guestuser2)
- VMware Identity Manager
- Lightweight Directory Access Protocol (LDAP)
- Principal Identity (using certificate or Jason Web Token)

To add user role assignments in projects, the following two methods are available:

Method 1: Add User Role Assignments from the User Management Page

The **User Management** page is available only to the enterprise administrator. A project administrator cannot use this page even if an enterprise administrator has granted permissions to the **Project Admin** role to do user role assignments in projects.

Method 2: Add User Role Assignments from the Manage Projects Page

The **Manage Projects** page is available to both the enterprise administrator and the project administrator. However, a project administrator can add user roles from this page only when an enterprise administrator has granted permissions to the **Project Admin** role to do user role assignments.

Read the following topics to know more.

- Add Role Assignments for a Project from the User Management Page

  Use the instructions in this documentation to add user role assignments for a project from the **User Management** page.

- Add Role Assignments for a Project from the Manage Projects Page

  Use the instructions in this documentation to add user role assignments for a project from the **Manage Projects** page.

Add Role Assignments for a Project from the User Management Page

Use the instructions in this documentation to add user role assignments for a project from the **User Management** page.

The **User Management** page is available only to the enterprise administrator. A project administrator cannot use this page even if an enterprise administrator has granted permissions to the **Project Admin** role to do user role assignments in projects.

The following procedure explains the steps for adding role assignments in projects for local user accounts and LDAP user accounts. The steps to add role assignments for vIDM user accounts remains the same, and therefore not covered in this procedure.

Prerequisites

User accounts are created. For example:

- Local user accounts are added in the system and they are activated.
- NSX Manager is configured to authenticate users from any of these identity management providers:
  - VMware Identity Manager (vIDM)
  - LDAP-based directory service, for example, Active Directory.

Procedure

1. From your browser, log in to an NSX Manager at `https://`*`nsx-manager-ip-address`*.
2. Navigate to **System > User Management**.

   The **User Role Assignment** tab is displayed.

3 To add role assignments in projects for a local user account, do these steps:

a Next to the local user account name, click the  icon, and then click **Edit**.

b Click the link under the **Roles** column.

The **Set Roles/Scope** dialog box opens.

c Click **Add Role**, and then select any one of these roles to assign to the local user:

- Project Admin
- Network Admin
- Security Admin
- Network Operator
- Security Operator

For example, let us assume that you have selected the **Project Admin** role.

d Under the **Scope** column, click **Set**, and then select one or more projects.

e Click **Add**, and then click **Apply**.

f Click **Apply** again to close the **Set Roles/Scope** dialog box.

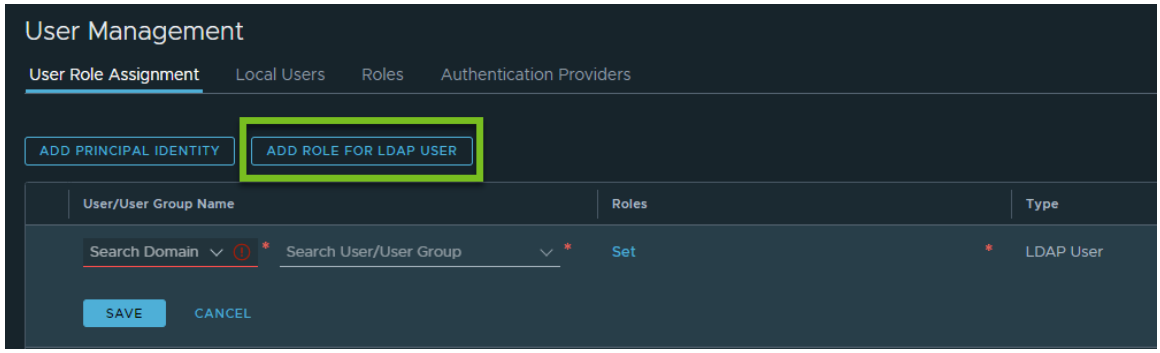
g Click **Save** to save the role assignment.

4 To add role assignments in projects for an LDAP user account, do these steps:

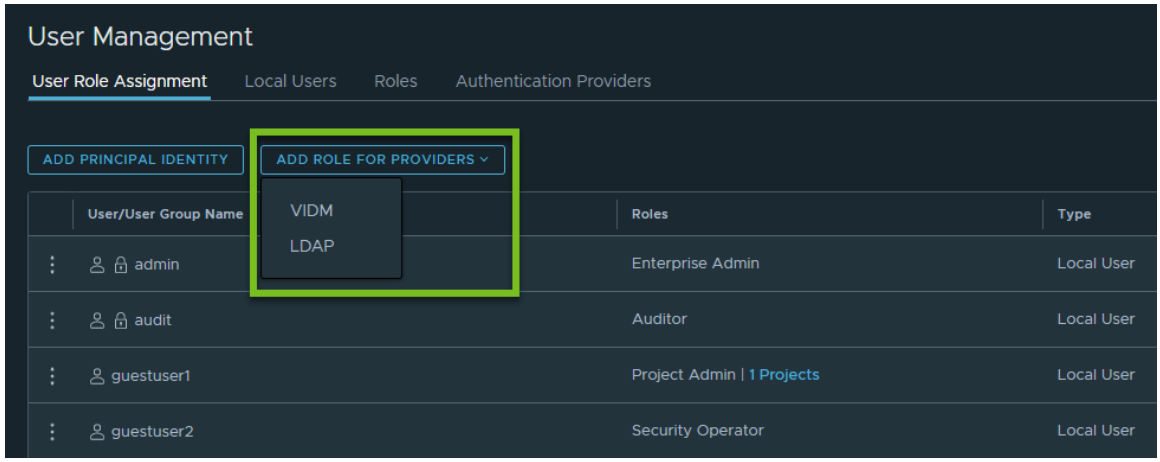
a Ensure that you are in the **User Role Assignment** tab.

b Click **Add Role for LDAP User**.

For example:



If you have configured NSX Manager to authenticate users from both LDAP and vIDM providers, the button caption is as shown in the following screen capture.



c Select a domain from the drop-down menu.

d Enter the first few characters of the user or group name.

System displays a list of matching user or group names. Select a user or a group name from the list.

e Under the **Roles** column, click **Set**.

The **Set Roles/Scope** dialog box opens.

f Click **Add Role** and follow the same process, as explained in steps 3(c) though 3(e), to assign roles to the LDAP user in the project.

Results

Project users can now log in to NSX Manager with their login credentials. Observe that the project switcher drop-down menu shows only those project names that the logged-in user has access to.

Add Role Assignments for a Project from the Manage Projects Page

Use the instructions in this documentation to add user role assignments for a project from the **Manage Projects** page.

The **Manage Projects** page is available to both the enterprise administrator and the project administrator. However, a project administrator can add user roles from this page only when an enterprise administrator has granted permissions to the **Project Admin** role to do user role assignments.

Prerequisites

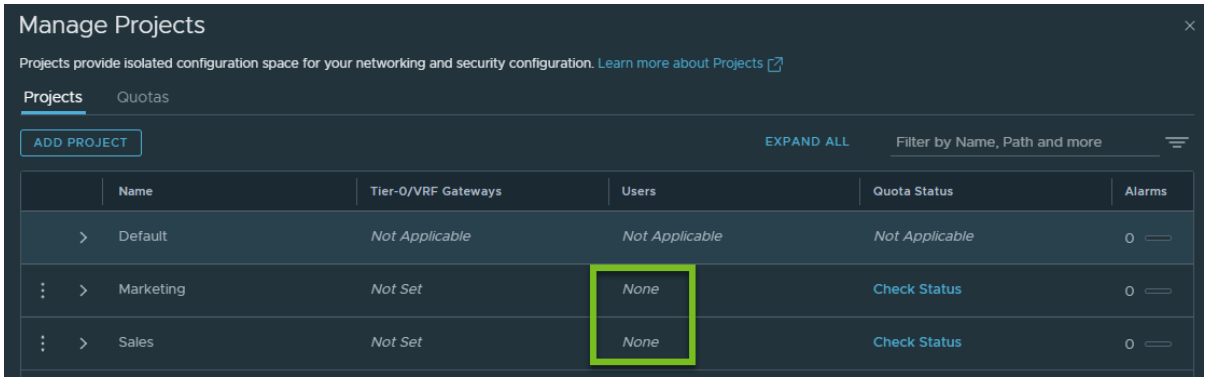
User accounts are created. For example:

- Local user accounts are added in the system and they are activated.
- NSX Manager is configured to authenticate users from any of these identity management providers:
  - VMware Identity Manager (vIDM)
  - LDAP-based directory service, for example, Active Directory.

Procedure

1. From your browser, log in to an NSX Manager at `https://nsx-manager-ip-address`.
2. Expand the project switcher drop-down menu from the application bar at the top, and then click **Manage**.

   For example, the following screen capture shows two projects on the **Manage Projects** page: Sales and Marketing. No users are currently set for these projects.



3. Next to the project name, click the icon, and then click **Edit**.
4. Under the **Users** column, click **Set**.

5 Click **Add Role Assignment**, and then do the following steps:

   a Select any one of these options:

      - Local User
      - LDAP User
      - vIDM User

   b Enter a user name or a group name.

      Note Ensure that the user or group name is entered accurately. Currently, the system does not display the list of matching users or groups when you type the first few characters in the text box. This behavior is a known limitation.

   c Under the **Role** column, click **Set**, and then click **Add Role**.

   d From the **Select Role** drop-down menu, select a role.

      For example, select the **Project Admin** role. Observe that the scope of this role is to set to `Project`. The scope of the role is non-editable.

   e Click **Add**, and then click **Apply**.

   f Click **Save** to save the role assignment.

6 (Optional) To add more user role assignments in the project, click **Add Role Assignment** again and repeat step 5.

Results

Project users can now log in to NSX Manager with their login credentials. Observe that the project switcher drop-down menu shows only those project names that the logged-in user has access to.

Add Quota to a Project

After adding a project in your org, you can add one or multiple quotas to define a maximum limit for the number of objects of a specific type that users can create in the project.

If you do not apply a quota to a project, system-defined maximum limits apply to the number of objects that users can create in that project.

Important The maximum limits that you define in the quota apply only to the creation of objects of a specific type in a project. Objects created in a project are owned by that project. Quotas do not put any constraints on the number of objects that an enterprise administrator can share with a project through resource shares. For example, if a project has a quota, which defines a limit of 10 groups, it does not restrict the enterprise administrator from sharing 20 groups with this project through a resource share.

To learn about resource shares, see Resource Sharing.

NSX has a validated maximum scale for a deployment. To know the maximum limits for NSX objects, see the VMware Configuration Maximums portal. The object limits mentioned on this portal apply to the entire system, meaning that the limits apply to objects that are both in and out of the projects. For example, the portal mentions that the maximum limit for segments is 10,000. This limit refers to the maximum number of segments in the system. Unless specified on the portal, there is no separate maximum limit for objects under each project in an NSX deployment.

By limiting the number of objects that users can create in a project, an enterprise administrator can govern the project resource consumption. For example:

- Ensure that one project does not take too much resources.
- Prevent project users from creating some objects by setting the quota to 0.
- Offer a different level of service between projects.

System allows you to reduce the quota limits of an object even below the current consumption.

For example, assume that for a project, a quota of 20 segments is defined, and 15 segments have been created. Later, the enterprise administrator decides to reduce the quota to 10 segments for this project. In this case, the consumption of segments has reached 150% of quota. However, existing segments remain in the project until they are deleted. Project users cannot create new segments inside the project until the segment count goes below 10.

Prerequisites

You must be assigned the **Enterprise Admin** role.

Procedure

1. From your browser, log in to an NSX Manager at `https://`*`nsx-manager-ip-address`*.
2. Expand the project switcher drop-down menu from the application bar at the top, and then click **Manage**.
3. Click the **Quotas** tab, and then click **Add Quota**.
4. Configure the limits for one or multiple objects in the quota.
   a. Enter a name for the quota.
   b. In the **Limits** column, click **Set**.

      The **View Limits** window opens. By default, you are in the **General** tab. On this tab, you can set limits on objects such as, redirection policy, redirection rule, and so on.
   c. Click the tab that contains the objects that you want to set limits on.
   d. Click **Add Limit**.
   e. From the **Objects** drop-down menu, select an object.

      For example, let us click the **Inventory** tab, and then select the **Group** object.

f In the **Limit** box, enter an integer, and then click **Save**.

For example, let us enter **10**.

Note To prohibit users from creating an object in the project, set the limit to **0** for that object.

g (Optional) Add limits for other objects, as required.

h After adding limits for all objects, click **Apply**.

5 From the **Apply To** drop-down menu, select a project to apply this quota.

A quota can be applied to multiple projects.

6 Click **Save**.

Example: Quota

In the example that we took earlier in this procedure, we defined a limit of 10 for the group object. Let us assume that this limit is applied to two projects: Project-1 and Project-2.

This quota puts a limit on creating a maximum of 10 groups in both the projects. After this limit is reached, you cannot add additional groups in your project. However, you can edit the limit later or delete it, if required.

When the number of user-created objects in a project reaches the limit that is set for an object, an alarm is triggered in the system. Currently, the alarm is visible only to the enterprise administrator and not to the project administrator. To view the alarm, navigate to **Home > Alarms** and check the details on the **Alarms** page. Currently, the system triggers an alarm only when the creation count limit for an object has reached in a project (that is, 100% quota is used).

Note For the `Creation Count Limit` event, the threshold value is currently not configurable. System-defined threshold value is used, by default.

What to do next

After a quota is applied to the project, a project administrator can log in to NSX Manager and monitor the quota status periodically for all the managed projects.

To know more about monitoring quota status, see Monitor Quota Status of a Project.

Monitor Quota Status of a Project

Project users can periodically check the quotas status of their projects to know the percentage of quota that is used, view any alarms that might have been raised, and decide whether any troubleshooting actions are required to resolve those alarms.

Procedure

1 From your browser, log in to an NSX Manager at `https://nsx-manager-ip-address`.

2 Expand the project switcher drop-down menu from the application bar at the top, and then click **Manage**.

3 In the **Quota Status** column, click **Check Status**.

- If the status changes to `OK`, it means that the number of user-created objects in the project is within the limit.
- If the status changes to `Error`, it means that the quota has reached (that is, 100% of quota is used).
- If the status changes to `Not Set`, it means that no quota is defined for the project.

Use the information that is displayed in the **Alarms** column to know whether any alarms are raised. Alternatively, you can navigate to **Home > Alarms** and view the active alarms, if any.

When the quota limit is reached, a `Creation Count Limit Reached` event is displayed on the **Alarms** page. In NSX 4.1, the alarm due to this event is visible only to the enterprise administrator, and not to the project administrator. View the details of this event to learn about the recommended actions.

Note For the `Creation Count Limit` event, the threshold value is currently not configurable. System-defined threshold value is used, by default.

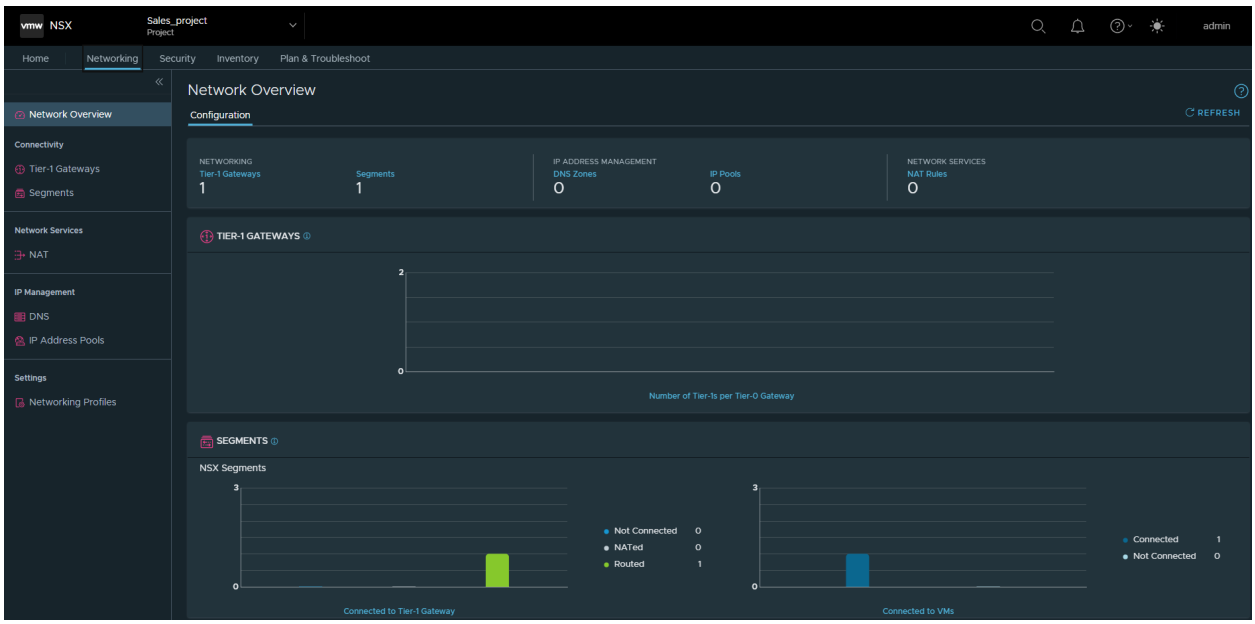
4 To know more about the quota details of a project, click the value (`OK` or `Error)` in the **Quota Status** column.

The **Quota Status** window displays quota details, such as the limits set for the various object types in the quota, the quota name, the percentage of quota used, and other details.

Networking Within a Project

This topic provides a brief overview of the NSX networking features that are currently supported inside projects.

The following screen capture shows the features that are available on the **Networking** page in the **Sales_project** view.



Managing Tier-1 Gateways in a Project

- Tier-1 gateways and segments inside a project can connect to tier-0 gateways that are assigned to the project.
- Tier-1 service routers and DHCP are provisioned on edge clusters that are assigned to the project.
- Tier-1 service routers that are provisioned on the edge clusters will benefit from the configuration that an enterprise administrator has done on the edge cluster, for example, configuration of the failure domain.
- Tier-1 service interfaces can connect to networks that are created within the projects and to networks that are shared with the projects.
- Features such as NAT, Gateway Firewall can be configured on the tier-1 gateways that are created inside a project. Tier-1 gateways that are configured in the default space cannot be allocated to projects.

Managing Networks in a Project

- Currently, overlay segments are supported inside projects.
- Segments inside a project can use the default transport zone of the system.
- DHCP profiles can be created on edge clusters that are allocated to the project.
- Segments can use the DHCP profiles that are shared with the project.
- To provide distributed security to VMs inside a project, these VMs must be connected to the segments inside a project.

Groups in a Project

NSX supports the creation of groups when setting up multi-tenancy in your environment.

Default Group in a Project

NSX creates a default group of type **Generic** for every project that is added in your environment. This default group represents the project itself.

By default, the system adds all the segments that users have created inside a project to the project's default group. Only those VMs that are attached to the segments in the project are members of the default group. The default group helps you to restrict the scope of the firewall rules to a particular project.

The following naming convention is used to identify the default group in a project:

`ORG-default-PROJECT-<Project_Name>-default`

User-created Groups in a Project

When you navigate to the **Set Members** page in the UI to select members for a group under a project, the system displays only those objects that are owned by your project or shared with your project.

The following NSX objects are supported for adding statically to a group definition inside a project:

- Segments
- Segment Ports
- VIFs
- Virtual Machines
- Groups

The following NSX objects are supported for adding to a dynamic group membership criteria inside a project:

- Virtual Machine
- Segment
- Segment Port
- Group

The following points apply to groups that are created inside the **Default** view (default space) of your NSX environment:

- Dynamic group membership evaluates to all the VMs in the NSX system, including the VMs in the projects.

For example, if a group in the default space includes VMs with the **Web** tag, the effective members of this group will include VMs with the **Web** tag both within and outside the projects.

Note Tags are metadata of the VMs in NSX. If a VM is attached to a segment in a project, both the project administrator and the enterprise administrator can assign tags to the VM. For example, let us say that in the default space, the enterprise administrator has assigned a tag to a VM that is inside project 1. The project administrator of project 1 can modify the tag that is assigned to this VM because the VM is seen in the project 1 inventory.

- For groups with static membership in the default space, you can add workload VMs that are connected to a project either by explicitly referring to the VMs (Members > Virtual Machines) or by adding the project default groups (`ORG-default-PROJECT-<Project_Name>-default`).

  Other objects created under a project cannot be added as static members to groups inside the default space.

Adding Group to a Project

The UI workflow for adding a group in a project remains the same as it currently exists for adding groups in the **Default** view (default space) of your NSX deployment.

The only difference is that in the UI, you must first select a project from the project switcher drop-down menu on the application bar at the top, and then navigate to **Inventory > Groups** for adding groups in that selected project.

Distributed Firewall Policies in a Project

NSX supports the creation of distributed firewall policies under projects.

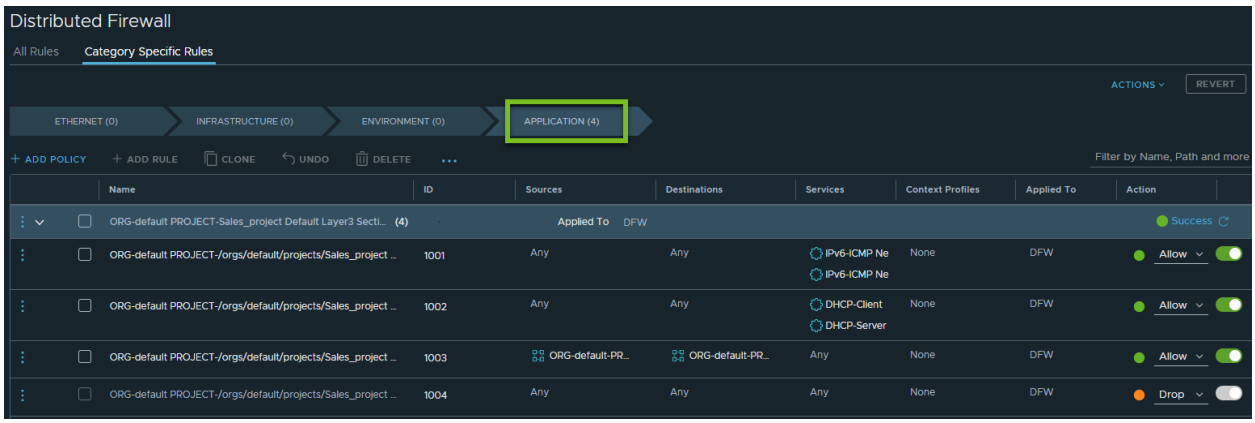
Default DFW Rules in a Project

For each project that is added in the NSX environment, the system creates a default DFW policy in the project. The default policy appears at the bottom of the policy list in the Application firewall category. The default policy defines the behavior for VMs within the project if no other rules are encountered.

The following naming convention is used to identify the default DFW policy in a project:

`ORG-default-PROJECT-<Project_Name>-Default Layer 3 section`

For example, the following screen capture shows the default DFW policy rules in the Sales_project.



As seen in this screen capture, the default policy is applied to DFW, and it contains the following four firewall rules:

- Rule 1001 allows IPv6-ICMP traffic.
- Rule 1002 allows communication to the DHCP client and server.
- Rule 1003 allows communication between workload VMs within the project.
- Rule 1004 drops all other communication.

The default DFW policy ensures that VMs within a project can only reach other VMs in the same project, including the DHCP server. Communication with VMs outside the project is blocked. The VMs that are connected to the segments inside the project cannot ping their default gateway, by default. If such a communication is required, you must add new rules or modify existing rules in the default DFW policy.

User-created DFW Rules in a Project

The Infrastructure, Environment, and Application DFW firewall categories are supported for projects within the org. Inside each firewall category, the DFW policies in the default space have the highest precedence followed by the project-specific policies.

The DFW rules in the default space can extend to a project.

Note The DFW rules in the default space apply to every VM in the NSX deployment, including the VMs in the projects. However, you can restrict the scope of the rules in the default space, by selecting the **Groups** option from the **Applied To** setting in the UI.

For instance, you can choose to apply the rules to the project default group (`ORG-default-PROJECT-<Project_Name>`). The project default group contains only the workload VMs of a project.

DFW rules inside a project have access to groups that are created in the project and to groups that are shared with the project. Groups that are shared with the projects can be used only in the **Source** or **Destination** fields of the firewall rule, and not in the **Applied To** field of the firewall rule.

The firewall rules in a project apply only to the VMs in the project, that is, VMs that are connected to the segments in the project. The firewall rules within a project, including the Any-Any rules that are applied to DFW, do not impact workloads outside the project.

Adding DFW Policy in a Project

The UI workflow for adding a DFW policy in a project remains the same as it currently exists for adding policies in the **Default** view (default space) of your NSX deployment.

The only difference is that in the UI, you must first select a project from the project switcher drop-down menu on the application bar at the top, and then navigate to **Security > Distributed Firewall** for adding DFW policies in that selected project.

Monitoring and Troubleshooting within Projects

NSX supports the monitoring of network traffic within Projects for debugging or troubleshooting purposes. You can also view alarms raised for a Project. The Enterprise Admin can view alarms within the default space and across all projects while the Project Admin can only view project-specifc alarms.

Traceflow

The Enterprise Admin and the Project Admin can use the Traceflow tool to inspect the path of a packet. While the Enterprise Admin can run a trace across the default space and across all projects, the Project Admin can perform a trace only within a project.

To run a trace across projects as an Enterprise Admin, ensure that you select **All Projects** from the project switcher menu. The UI workflow for performing a Traceflow in a project remains the same as it currently exists for the Default view (default space) of your NSX deployment.

To perform a trace as a Project Admin, navigate to **Plan & Troubleshoot > Traffic Analysis > Traceflow > Get Started** in your project space. Specify the source and destination information according to the traffic type and click **Trace**. The output includes a table listing the observed packets. Details of entities that belong to the default space and encountered during the trace, are hidden from view in the Traceflow output for the project.

For more information on Traceflow, see Perform a Traceflow.

Port Mirroring

Only the Enterprise Admin can configure port mirroring from the default space for VMs in a project. This feature is currently unavailable to the Project Admin and only the Enterprise Admin can view the results in the default space.

To configure port mirroring for a project, create appropriate Groups in the default space and add your tagged VMs from the project to these groups.

Navigate to **Plan & Troubleshoot > Port Mirroring**, select the session type, and set the groups as **Source** and **Destination**.

For more information on Port Mirroring, see Add a Port Mirroring Session.

IPFIX

Only the Enterprise Admin can configure IPFIX from the default space for VMs in a project. This feature is currently unavailable to the Project Admin and only the Enterprise Admin can view the results in the default space.

To configure IPFIX for a project, create appropriate Groups in the default space and add your tagged VMs from the project to these groups.

Navigate to **Plan & Troubleshoot > IPFIX**, select the appropriate group under **Applied To** for your IPFIX profile.

For more information on IPFIX, see Chapter 21 Network Monitoring.

Features Available for Consumption Under Projects

In a multi-tenant NSX environment, project administrators can configure a subset of NSX features under projects.

This provides flexibility to the tenants while allowing enterprise administrators and other administrators of the default space to control the overall system life cycle and connectivity.

Table 19-1. Features Available Under Projects

NSX Feature	Available Under Projects	Notes
System		
Edge Clusters	Assigned during project creation	During project creation, the enterprise administrator assigns an edge cluster from the default space to a project.
System Life cycle	Managed by the enterprise administrator	Platform-wide operations such as install, upgrade, and backup are managed by the enterprise administrator.
Networking		
Tier-0 or Tier-0 VRF Gateway	Assigned during project creation	During project creation, the enterprise administrator assigns a tier-0 gateway or a tier-0 VRF gateway from the default space to a project. The services on the tier-0 gateway are managed by the enterprise administrator from the default space.
Dynamic Routing (BGP/OSPF)	Managed by the enterprise administrator	Dynamic routing is configured on the tier-0 gateway or tier-0 VRF gateway by the enterprise administrator.
EVPN	Managed by the enterprise administrator	EVPN is configured on the tier-0 gateway or tier-0 VRF gateway by the enterprise administrator.
Tier-1 Gateway	Yes	
Static Routing	Yes	Static routing on the tier-1 gateway is configured by the project administrator.

Table 19-1. Features Available Under Projects (continued)

NSX Feature	Available Under Projects	Notes
Overlay Segments	Yes	
VLAN Segments	No	
Segment Profiles ■ SpoofGuard ■ IP Discovery ■ MAC Discovery ■ Segment Security ■ QoS	Yes	
L2 Bridge	No	
L2 VPN	No	
L3 VPN	No	
NAT	Yes	
Load Balancer	No	
DNS Forwarder ■ DNS Services ■ DNS Zones	Yes	
IP Address Pools/IP Address Blocks	Yes	
IPv6 Profiles (DAD/ND)	Yes	
Gateway QoS Profiles	Yes	
DHCP and DHCP Relay	Yes	
Security		
Distributed Firewall	Yes	Only for VMs that are connected to segments in a project. Firewall rules that are managed by the enterprise administrator in the default space have the highest precedence followed by the project policies. DFW policies with Antrea groups in the Sources, or Destinations, or Applied To of the firewall rules are not supported under projects.
Exclusion List	Managed by the enterprise administrator	Exclusion list excludes a VM from all the firewall application rules.

Table 19-1. Features Available Under Projects (continued)

NSX Feature	Available Under Projects	Notes
Gateway Firewall	Yes	The enterprise administrator and project administrator can manage the gateway firewall rules on the tier-1 gateways of the project only within the context of a project. The project administrator can delete or modify the gateway firewall rules in the project that are created by the enterprise administrator.
Identity Firewall	No	Identity firewall is not available under projects. Identity firewall rules can be configured only in the default space and these rules can be applied to VMs inside projects.
Distributed IDS/IPS Gateway IDS/IPS Malware Prevention TLS Decryption FQDN Filtering	No	
TLS Inspection FQDN Analysis URL Filtering	No	
Firewall Profiles ■ Session Timer ■ Flood Protection ■ DNS Security	Yes	
Inventory		
Services	Yes	
Groups (static and dynamic memberships)	Yes	Kubernetes member types are not available under a project for creating dynamic membership criteria.
Antrea Groups	No	
Context Profiles/L7 Access Profiles	Yes	
Tags	Yes	
Virtual Machines (visibility/tagging)	Yes	Only for VMs that are connected to segments in a project.
Container Clusters	No	Kubernetes resources in Antrea Kubernetes clusters that are registered to NSX are not exposed to the project inventory.
Plan and Troubleshoot		

Table 19-1. Features Available Under Projects (continued)

NSX Feature	Available Under Projects	Notes
Traceflow	Yes	Traceflow can only use VMs and ports that are a part of the project. If the destination is an IP that routes to another project, these details are hidden in the Traceflow output.
Antrea Traceflow	No	
Live Traffic Analysis	No	
IPFIX	Managed by the enterprise administrator	IPFIX is managed centrally by the enterprise administrator in the default space.
Port Mirroring	Managed by the enterprise administrator	Port mirroring is configured centrally by the enterprise administrator in the default space.
NSX Intelligence	No	NSX Intelligence features are not exposed to the project administrator. Only the NSX Enterprise Admin user role has full access to all the NSX Intelligence features. The NSX Intelligence features (network flow visualization, micro-segmentation recommendations, and suspicious traffic analysis) do not work with isolated security and networking configurations, such as projects. These features work with all the network traffic flow data within your entire on-premises NSX environment. If multi-tenancy is used, that is, if projects are defined in your NSX environment, NSX Intelligence displays all the NSX objects that are in the default space and all the NSX objects from all the projects, whether you are using the **Default** view or the **All Projects** view.

System Monitoring

20

You can monitor the health and performance of the NSX environment.

This chapter includes the following topics:

- Monitor NSX Edge Nodes and Gateways
- APIs to Fetch Time-Series Metrics
- Dynamic Plugins
- Working with Events and Alarms
- Registering Notification Watchers
- Using Aria Operations for Logs or Splunk for System Monitoring
- Using Aria Operations Manager for System Monitoring
- Using Aria Operations for Networks for System Monitoring

Monitor NSX Edge Nodes and Gateways

NSX application supports collection and storage of data up to one year. With this feature, along with the point-in-time data, you can also view time series metrics, such as CPU usage memory, disk usage, packets per second, bytes per second, and packet drop rate, for edge nodes and gateways. You can view the metrics by invoking time series metrics APIs. A few of these metrics are also available on NSX Manager interface along with the point-in-time data.

Starting with NSX 4.0.1.1, you can also view time series metrics for VPN. The VPN time series metrics is not available on NSX Manager. You can use time series APIs to view data such as policy based VPN statistics and route based VPN statistics, which includes metadata about number of tunnels configured and how many tunnels are up or down at a certain time.

Using time series metrics, you can monitor the trend in key performance indicators, detect anomalies, perform before and after analysis, and get the historical context which can help in troubleshooting. Based on your role, you can view metrics of only those objects for which you have the authorization.

You can use the metrics APIs to fetch the time series metrics. For information about using the APIs, see APIs to Fetch Time-Series Metrics.

On the NSX Manager interface, the time series metrics is displayed as charts with respective keys and time filter on last one hour, last 24 hours, last one week, last one month, and last one year. You can view the following information for edge nodes:

- CPU cores allocated
- Alarms
  - Edge node - Overall alarm count
  - CPU - Alarms for Datapath CPU and Service CPUs
  - Disk - Overall disk alarms and alarms for each partition
  - Memory - Overall memory alarms and alarms for each memory pool
- Uptime
- CPU
  - System Load - You can view current load average and average system load trend.
  - Datapath CPU - Number of datapath CPU cores and their usage details, which include the average usage of all cores and the highest usage among cores. You can view current utilization, highest utilization for a single core, and average usage trend. Starting with NSX 4.0.1.1, you can also view the average number of packets processed across all available datapath cores for a given edge transport node and which core has processed the highest number of packets.
  - Service CPU - Number of service CPU cores and their usage details, which include the average usage of all cores and highest usage among cores. You can view current utilization, highest utilization for a single core, and average usage trend.
- Memory
  - System Memory - Utilization of RAM on edge node. You can view current utilization and average usage trend.
  - Datapath Memory - Includes heap memory, memory pool, and resident memory. You can view current usage, current pool usage, and average usage trend.
  - Memory Pools - List of all memory pools along with their description and usage values except for QAT memory pool (of Bare Metal Edge) whose usage is always around 100%. The memory pools are:

Name	Description
jumbo_mbuf_pool	Packet Pool for Jumbo Frame Used by Ipsec Crypto Device
common_mbuf_pool	Datapath Common Packet Pool
sp_pktmbuf_pool	Datapath Slowpath Packet Pool
fw_mon_msg	Stateful Service Sync Message Pool
vxstt4_frag_q	Vxstt Fragment Pool for Reassembly
pfstatepl3	Stateful Service State Pool

Name	Description
pffqdnippl	Stateful Service FQDN to IP Map Pool
pffqdnsyncpl	Stateful Service FQDN SYNC Pool
pffqdndnpl	Stateful Service FQDN Internal Pool
pfdnsdnpl	Stateful Service FQDN Internal Pool
pfpktpl3	Stateful Service Fragmented Packet Pool
pfsyncmbufpl3	Stateful Service SYNC Pool
pf_fp_rule_node	Stateful Service Rule Node Pool
pf_fp_root_rule_node	Stateful Service Root Rule Node Pool
pf_tb_root_rule_node	Stateful Service Fastpath Root Table Node Pool
pfa_intattr_pl3	Stateful Service Integer Attribute Pool
pfa_attrconn_pl3	Stateful Service Attribute Connection Pool
pfa_ctx_pl3	Stateful Service Context Pool
pfa_key_ace_pl3	Stateful Service Integer Attribute Key Pool
pfa_value_ace_pl3	Stateful Service Integer Attribute Value Pool
lb_pkt_pl3	Load Balancer Temp Packet Cache Pool

- Disk
  - Total disk usage for all ext4 disk partitions and the list of RAM disk and disk partitions. Also, the available free space of each partition. You can view overall GB used, read latency, write latency, current usage, average usage trend by partition, and average disk latency trend.
- Transport node status
- Network interface
  - To view statistics of any network interface, click the graphic icon. You can view cumulative statistics like total packets and dropped packets and time series metrics like network utilization trend and dropped packets trend. Starting with NSX 4.0.1.1, you can also view Error chart for a gateway.
- NAT rules statistics

Prerequisites

Ensure that you have deployed NSX Application Platform. For more details about deploying NSX Application Platform, see the *Deploying and Managing NSX Application Platform* guide. The Metrics feature is enabled by default when you deploy NSX Application Platform.

If you view statistics without deploying NSX Application Platform or without enabling the Metrics feature, the system displays a message to deploy NSX Application Platform and to enable the Metrics feature.

Procedure

1 To view metrics of an edge node, perform the following steps:

   a With admin privileges, log in to NSX Manager.

   b Select **System > Fabric > Nodes**.

   c Click the **Edge Transport Nodes** tab.

   d Click the edge node that you want to monitor.

   e Click the **Monitor** tab.

   Usage information for CPU, memory, and disk is displayed, as well as the node status, network interfaces, and NAT rule statistics.

2 To view gateway or interface metrics, perform the following steps:

   a With admin privileges, log in to NSX Manager.

   b Select **Networking > Network Topology**.

   c Click the required gateway. The system displays a gateway details window.

   d Under **Statistics**, select the linked gateway or the interface link for which you want to view the metrics.

APIs to Fetch Time-Series Metrics

You can use the metrics APIs to fetch the time series metrics.

These APIs that can take multiple intent paths as input for a specific resource type, such as edge or gateway, and return the corresponding metrics. For complete information about how to invoke the time series metrics APIs, see *NSX Intelligence & NSX Application Platform API Guide*.

The following list contains terms and their descriptions related to metrics APIs.

Term	Description
Resource type	Any entity, such as edge node or gateway, for which metrics is available.
Resource ID	Intent path or an identifier for a resource type.

Term	Description
Object	Sub-entity of a particular resource type. For example, firewall rule within a ruleset.
Metric key	Unique key assigned to a metric of any resource. Following metric keys are examples of keys available for edge nodes and NSX Manager: ■ disk.avg_used_percent: Percent of blocks in use over total. ■ system.avg_load: The average system load over the last 1 minutes. ■ system.avg_mem_available_percent: Percent of available memory over total memory. Following metric keys are examples of keys available for gateway firewalls: ■ edge_fw.avg_drop_reason_alg: Average rate of firewall drop due to Application Layer Gateway. ■ edge_fw.avg_drop_reason_connection_limit: Average rate of drop per second due to connection-limit. You can use the following API to fetch metric keys of any required resource. `GET https://<manager>/napp/api/v1/metrics/key-info?resource_type=<resource_type>`

Perform the following steps to fetch metrics for a required resource. Note that metric keys used in examples are for the purpose of sample only. Use the appropriate API to fetch actual keys and descriptions.

Step 1: Select the resource type for which you want to fetch the time-series metrics

You can fetch time series metrics for edge nodes and gateways. Select the resource type from the following list.

- PolicyEdgeNode
- ClusterNode
- Tier0Interface
- Tier1Interface
- Tier0
- Tier1

  Starting with NSX 4.0.1.1, you can also view time series metrics for VPN by using the following resource types.

- PolicyBasedIPSecVpnSession
- RouteBasedIPSecVpnSession

Step 2: Get supported metric keys and their information for a resource type

Use the following API to get supported metric keys for a resource type. The API takes `resource_type` as a query parameter and list all metrics that are available for that resource. This API also provides description and units for each metric.

URI Path:

```
GET https://<manager>/napp/api/v1/metrics/key-info?resource_type=<resource_type>
```

Example:

```
GET https://<manager>/napp/api/v1/metrics/key-info?resource_type=PolicyEdgeNode
```

Response:

```
{
  "results":[
        {
            "metric_key ":"edge.cpu_usage",
            "metric_unit":"PERCENT",
            "description":"Edge Cpu usage percentage"
        },
     {
            "metric_key":"edge.pnic_avg_rx_packets",
            "metric_unit":"PER_SECOND"
            "description":"Average Rx packets per second"
        }
  ]
}

}
```

Step 3: Get object information for a resource and a metric key

Use the following API to get the object information of a metric key for a required resource type.

URI Path:

```
POST https://<manager>/napp/api/v1/metrics/object-info
```

Example Request:

```
{
    "resource_type":"PolicyEdgeNode",
    "resource_ids": [
        "/infra/sites/default/enforcement-points/default/edge-clusters/57d2c653-4d63-48d8-
b188-40b4e45a9bc8/edge-nodes/2ed9af04-21c9-11e9-be65-000c2902dff7",
        "/infra/sites/default/enforcement-points/default/edge-clusters/57d2c653-4d63-48d8-
b188-40b4e45a9bc8/edge-nodes/1349af04-21c9-11e9-be65-000c2902d0000"
    ],
    "keys": [ "edge_cores.cpu_usage", "edge.pnic_avg_rx_packets"],
```

```
    "start_time":1603971420
    "end_time": 1603973420
    "granularity": "5M",
    "max_num_data_points": 10
}
```

Example Response:

```
{
    "start_time": 1603971420,
    "end_time": 1603973420,
    "resource_type": "PolicyEdgeNode",

    "results": [
        {
            "resource_id": "/infra/sites/default/enforcement-points/default/edge-clusters/
57d2c653-4d63-48d8-b188-40b4e45a9bc8/edge-nodes/2ed9af04-21c9-11e9-be65-000c2902dff7",
            "key_results": [
                {
                    "key": "edge.pnic_avg_rx_packets",
                    "unit": "PER_SECOND",
                },
                {
                    "key": "edge_core.cpu_usage",

                    "results": [
                        {
                            "object_id": "core1",
                            "node_id": "2ed9af04-21c9-11e9-be65-000c2902dff7",
                            "node_path": "/infra/default/edge/1",
                             node_name": "Edge1",

                        },
                        {
                            "object_id": "core2",
                            "node_id": "2ed9af04-21c9-11e9-be65-000c2902dff7",
                            "node_path": "/infra/default/edge/1",
                            "node_name": "Edge1",
                        }
                    ]
                }
            ]
        },
        {
            "resource_id": "/infra/sites/default/enforcement-points/default/edge-clusters/
57d2c653-4d63-48d8-b188-40b4e45a9bc8/edge-nodes/1349af04-21c9-11e9-be65-000c2902d0000",
            "key_results": [
                {
                    "key": "edge.pnic_avg_rx_packets",
                    "unit": "PER_SECOND",
                    "description": "Average packet per second (PPS) network utilization on
ingress on the the management interface(s)",

                }
```

```
            ]
        }
    ]
}
```

Step 4: Fetch metrics data

Use the following API to fetch metrics data.

URI Path:

```
POST https://<manager>/napp/api/v1/metrics/data
```

Example Request:

```
{
    "resource_type":"PolicyEdgeNode",
    "resource_ids": [
        "/infra/sites/default/enforcement-points/default/edge-clusters/57d2c653-4d63-48d8-
b188-40b4e45a9bc8/edge-nodes/2ed9af04-21c9-11e9-be65-000c2902dff7",
        "/infra/sites/default/enforcement-points/default/edge-clusters/57d2c653-4d63-48d8-
b188-40b4e45a9bc8/edge-nodes/1349af04-21c9-11e9-be65-000c2902d0000"
    ],
    "object_ids": ["core1" ,"core2"],
    "node_ids": [],
    "keys": [ "edge_cores.cpu_usage", "edge.pnic_avg_rx_packets"],
    "start_time":1603971420
    "end_time": 1603973420
    "granularity": "5M",
    "max_num_data_points": 10
}
```

Example Response:

```
{
    "start_time": 1603971420,  --> output start_time and end_time maybe different from input
start_time and end_time
    "end_time": 1603973420,
    "resource_type": "PolicyEdgeNode",

    "results": [
        {
            "path": "/infra/sites/default/enforcement-points/default/edge-clusters/
57d2c653-4d63-48d8-b188-40b4e45a9bc8/edge-nodes/2ed9af04-21c9-11e9-be65-000c2902dff7",
            "key_results": [
                {

                {
                    "key": "edge_core.cpu_usage",

                    "results": [
                        {
                            "object_id": "core1",
                            "node_id": "2ed9af04-21c9-11e9-be65-000c2902dff7",
```

```
                            "data": [
                            {
                                "time": 1603444589,
                                "value": 29
                            },
                            {
                                "time": 1603444489,
                                "value": 30
                            },
                            {
                                "time": 1603444389,
                                "value": 35
                            }
                        },
                          {
                            "object_id": "core2",
                            "object_description": "Some description",
                            "node_id": "2ed9af04-21c9-11e9-be65-000c2902dff7",
                            "node_name": "Edge1",
                            "node_path":"/infra/default/edge/1"
                            "data": [
                                {
                                    "time": 1603444589,
                                    "value": 29
                                },
                                {
                                    "time": 1603444489,
                                    "value": 30
                                },
                                {
                                    "time": 1603444389,
                                    "value": 35
                                }
                           }
                    ]
                }
            ]
        }
```


Dynamic Plugins

A dynamic plugin is a customized plugin that you can create for any supported transport node, such as ESXi host, to check the node's health.

A dynamic plugin can be installed at runtime. It performs the following functions:

- Write a system log for the affected transport node.
- Execute a command or CLI through the `run_command()` function.
- Read existing metrics.
- Export data to wavefront (only in a VMware Cloud environment).

You can create a dynamic plugin if you have expertise in managing NSX. A dynamic plugin can also be created by VMware Support. A created plugin must be submitted to a GIT repository for validation. Before the submission, the plugin must be reviewed and tested.

Security management

After the completion of plugin review and test, the plugin, its test result, and code changes are submitted for validation in a GIT repository. Use the following GIT repository details

- GitLab: `https://gitlab.eng.vmware.com/core-build/nsbu-sha-plugin`
- Product: **`nsx-sha-plugins`**
- Gitreview supporting: Performed by VMware team

After the plugin is validated and approved, it is committed to the GIT repository. All dynamic plugins are built and signed when a new build is created. Each plugin is packaged and signed separately. You can get the required plugin from the published files of the build.

When a signed plugin is uploaded to the management plane, the management plane uses a public key to verify the signature and confirm that this plugin is valid. After the plugin is validated, the files of the plugin are pushed to the destination hosts through the secured channels between the management plane and the Central Control Plane (CCP) and between CCP and hosts.

If a System Health Agent (SHA) instance is restarted, it gets the plugin files from the management plane again. Since all the files are published through secured channels and no temporary files are used, the risk that hackers can modify scripts is prevented.

Also, to prevent risks of harmful code, SHA uses RestrictedPython to check the plugin python script before executing the script.

Version management

A plugin might be based on a command or a tool that is not supported in later versions of NSX, so each custom plugin must define the supported NSX version in the `manifest.yml` file. The version should be a REGEX string for all the supported versions. SHA on the host-side checks the version of the custom-plugin and runs only the REGEX matched ones.

Recommended version management policies are:

- Define the supported NSX version of a major release.

  Considering that most of the commands and tools do not change between minor releases in the same major release, the following method is the suggested way to define the version for all minor releases.

  For example,

  ```
  version: ^2\.5\.[0-9.]+ <== The custom plugin supporting all NSX 2.5 releases
  ```

- When a new major release is published, all submitted dynamic plugins should be reviewed.
- The plugin writer must update the script when the related commands or tools change.

Install a dynamic plugin

The transport node or the edge node on which the plugin is installed must have a minimum of 30-MB memory space. Also, note that you can install only up to 10 plugins. Once the plugin count reaches 10, any further installation of a plugin will fail.

To install the plugin, perform the following tasks:

1 Create the dynamic plugin files in the GIT repository. For more information about the plugin files, see the section *Dynamic plugin files*.

2 Trigger the product build in the GIT repository to generate the zipped package of the dynamic plugin files and download the package.

3 Create the dynamic plugin by using the following API with the POST method.

```
https://<manager_ip>/api/v1/systemhealth/plugins
```

4 Upload the plugin zipped package to the management plane using the following API with the POST method. The management plane extracts the uploaded file and perform the required validation.

```
/systemhealth/plugins/<plugin-id>/files/<file-name>/data
```

Note The maximum size of the plugin zipped file is 500k.

5 Create a node group with the required transport nodes as members by using the following API with the POST method.

```
/<manager_ip>/api/v1/ns-groups
```

6 Apply the plugin profile to the node group by creating a new service config by using the following API. The service config framework sends the plugin content to the node group.

```
https://<manager_ip>/api/v1/service-configs
```

For more information about APIs, see the *NSX API Guide* documentation.

Get the plugin status

Once the dynamic plugin is running, it automatically uploads the status to the management plane through the existing message channel. The management plane aggregates the plugin status information and store it into the database. To get the status of all plugins on each node, use the following API with the GET method.

```
https://<manager_ip>/api/v1/systemhealth/plugins/status/<transport_node_id>
```

Request Example:

```
GET https://<manager_ip>/api/v1/systemhealth/plugins/status/a257b981-1a1c-4b95-
b16c-8646
```

Response Example:

```
{
 "result_count":1,
 "results": [
 {
 "id": "72e1bd4b-6df6-42d0-9c59-a1c31312c9f1",
 "name": "health-check-compute",
 "status": "NORMAL",
 "detail": ""
 }
 ]
 }
```

Uninstall a dynamic plugin

To uninstall a plugin, remove the service config by using the following API.

`https://<manager_ip>/api/v1/service-configs/<service_config_id>`

Other APIs for managing the plugins

The following table lists APIs to manage dynamic plugins. For more information about APIs, see the *NSX API Guide* documentation.

Task	Method	API
Delete a plugin	DELETE	`/systemhealth/plugins/<plugin-id>`
Create a system health profile	POST	`/systemhealth/profiles`
Watch the plugin status	GET	`/systemhealth/plugins/status/<node-id>`

Enable the plugin		Enabling a plugin is a two-step process as follows:
		1 Use the following API to set the `enabled` property to `true` or `false`.
		`https://<manager_ip>/api/v1/systemhealth/profiles/`
		2 Use the following API to apply the SHA profile to the NS group.
		`https://<manager_ip>/api/v1/service-configs`

Change the plugin interval	POST	Changing the plugin interval is a two-step process as follows:
		1 Use the following API to set the `config` property.
		`https://<manager_ip>/api/v1/systemhealth/profiles/`
		2 Use the following API to apply the SHA profile to the NS group.
		`https://<manager_ip>/api/v1/service-configs`

Dynamic plugin files

A dynamic plugin comprises the following files:

- Install specification file

The install specification file, `manifest.yml`, contains the following information for System Health Agent:

- Plugin structure
- Constraints if any
- How to install and use the plugin
- Security restrictions for the health check script. For example, permissions the script has and files that the script can access.

The following table lists fields that are specified in a `manifest.yml` file.

Name	Description	Required/Optional	Example
classes	Specifies classes needed in the plugin script. The classes must be specified in the following format. '<module_name>.<class_name>'	Optional	`classes: ['datetime.datetime','datetime.date']`
modules	Specifies modules needed in the plugin script.	Optional	modules: ['random', 'math']

plugin	Specifies the plugin structure as follows:
	config: config file name
	script: script file name
		Required	`plugin:`
			`config: plugin.cfg.yml`
			`script: plugin.py`

version	Specifies the NSX versions on which this plugin can be installed.	Required	`version: '^3\.1\.[0-9.]+'`
node_type	Specifies NSX node types where this plugin can be installed. The available node types are: ■ nsx-esx ■ nsx-bms ■ nsx-edge	Required	`node_type: ['nsx-esx']`
metrics	Specifies metrics which can be consumed in the plugin script.	Optional	`metrics: ['nsx.host.host-metrics']`
precondition	Specifies precondition for the plugin. The available precondition is wavefront. **Note** This field is applicable only in a VMware Cloud (VMC) environment.	Optional	`precondition: ['wavefront']`

Do not use the following built-in modules:

- os
- subprocess
- sys
- multiprocessing
- importlib

The following table lists the interfaces that you must use in place of built-in functions of the respective modules. These interfaces are system provided. You can use them directly without specifying their module/class in the `manifest.yml` file.

Module	Built-in function	Substitute interface
datetime	datetime.date.strftime(self, fmt)	datetime_date_strftime(dt, fmt) :param dt: date instance :param fmt: format string
datetime	datetime.date.today()	datetime_date_today()
sys	sys.getrecursionlimit()	sys_getrecursionlimit()
sys	sys.getrefcount(object)	sys_getrefcount(object)
sys	sys.getsizeof(object, default)	sys_getsizeof(object, default)
sys	sys.maxsize	sys_maxsize
sys	sys.path	sys_path

Sample of a `manifest.yml` file.

```
# specify classes needed in plugin script
classes: ['datetime.datetime','datetime.date']
# specify modules needed in plugin script
modules: ['random', 'math']
# plugin structure
plugin:
 config: plugin.cfg.yml
 script: plugin.py
# specify nsx versions on which this plugin can be installed
version: '^3\.1\.[0-9.]+'
# specify nsx node type where this plugin can be installed
node_type: ['nsx-esx']
# specify metrics which can be consumed in plugin script
metrics: ['nsx.host.host-metrics']
# specify precondition for plugin
precondition: ['wavefront']
```

- Default profile file

The default profile file, `plugin.cfg.yml`, is used to configure plugin behavior, such as the execution frequency of the health check script. To change the default configurations, you can create a SHA profile for a specific dynamic plugin and apply it to transport node through NS group by using the management plane to CCP to NestDB channel.

The following table lists fields that are specified in a `plugin.cfg.yml` file.

Name	Description	Required/Optional	Example
CHECK_INTERVAL	Specifies the default interval in seconds for plugin script execution.	Required	CHECK_INTERVAL: 20
ENABLE	Specifies whether the plugin is enabled by default.	Required	ENABLE: true

Sample of a `plugin.cfg.yml` file.

```
# Required field - default interval (unit: second) between plugin script executions.
CHECK_INTERVAL: 20

# Required field - whether plugin is enabled by default
ENABLE: true

# Plugin user can add other fields as below if needed to control plugin script logic.
EXPORT_DATA: true
```

- Health check script

  A health check script file, `plugin.py`, contains a python script to check the health status of a transport node.

  The following table lists system-defined variables and functions that can be used and data that can be read in a `plugin.py` file.

Variable/Data/Function	Description	Type	Example
logger	Writes log information in syslog. The existing system-defined variable, logger, can be used directly in the plugin script. The output log is prefixed with the plugin name and id as shown in the following sample output. `2020-10-28T10:47:43Z nsx-sha: NSX 2101378 - [nsx@6876 comp="nsx-esx" subcomp="nsx-sha" username="root" level="INFO"] [name:hl-esx-002-04][id:a3eb14f1-d185-4bc7-bfaa-6cf888bbeb22] dynamic plugin - not export data`	Variable	`logger.info("this is a demo log")`
data_store	The existing system-defined dictionary that is used to fetch system-provided data. For example, profile, metric, and host_id.	Variable	`profile = data_store['profile']`
profile	Profile data is a dictionary parsed from the default profile (plugin.cfg.yml) or the effective SHA profile (user-applied through Manager API) that is read from data_store. It has the following format: {'ENABLE': True, 'CHECK_INTERVAL': 20, 'EXPORT_DATA': True}	Data	`profile = data_store['profile']`

metric	Metric is a dictionary with 'value' and 'timestamp' that is read from data_store. It has the following format:
	data_store['metrics'][<metric_name>]
	Where,
	The first key must be 'metrics'.
	The second key is an existing metric name.
		Data	`metric = data_store['metrics']['nsx.host.host-metrics']`
			`metric is: {`
			`'value':XXX, <== the collected data of the metric`
			`'timestamp': XXX <== timestamp of the data collected`
			`}`
			Note: The first run of plugin might not return a metric as currently a metric is collected asyncly with the plugin running, so the metric might not have been collected in the first run of the plugin.

host_id	Host_id is an instance of class uuid.UUID that is read from data_store.	Data	`host_id = data_store['host_id']`

Variable/Data/Function	Description	Type	Example
run_command	This function runs commands in a list format. It has the following format. `run_command(cmd, timeout=4)` Where, ■ *cmd*: Commands to be executed. Must be in the list format as in the example. ■ *timeout*: Timeout for waiting for the command result. Default timeout is 4s. Timeout should not be set larger than 20s. This function returns the command execution result.	Function	`cmd = ['nsxdp-cli', 'ipfix', 'settings', 'granular', 'get', '--dvs-alias', 'nsxvswitch', '--dvport=dafa09ca-33ed-4e04-ae3d-1c53305d5fe6']` `res = run_command(cmd)`
Exportdata	This function exports data to wavefront. Currently, a dynamic plugin supports export to wavefront only. **Note** This function is applicable only in a VMware Cloud (VMC) environment. It has the following format: `Exportdata: ExportData(data={}, exports=[], source=host_uuid)` Where, *data*: data to be exported; data should be in dictionary format as example. *exports*: destination list for export. In HL, only support wavefront in destination. It is required. *source*: source string for export. It is useful only for wavefront destination. It is optional, the default value is NSX host_uuid. The function does not return any value.	Function	`Exportdata(data={'esx.plugin.stats': {'stats': {'gc-esx-001': data}}}, exports=['wavefront'])`

Sample of a `plugin.py` file.

```
def report_test_data(edge_service_status):
    if edge_service_status == 'running':
        data = 2
    else:
        data = 3

    # examples on how to report data to wavefront.
    Exportdata(data={'esx.plugin.stats': {'stats': {'esx-dynamic-plugin-001': data}}},
exports=['wavefront'])


def run():
```

```
    # examples on how to write log.
    logger.debug("this is a debug message!")
    logger.info("this is a demo message!")

    # examples on how to use specified module in manifest. Take 'random' as an example.
    s_res = random.randint(1,10)
    logger.warn("random.randint(1,10)=%s", s_res)

    # examples on how to use specified class in manifest. Take 'datetime' and 'date' as an
example.
    logger.info('date.ctime(datetime.now()):{}'.format(date.ctime(datetime.now())))

    # examples on how to run cmd via interface run_command
    cmd = ['nsxdp-cli', 'ipfix', 'settings', 'granular', 'get', '--dvs-alias',
'nsxvswitch', '--dvport=dafa09ca-33ed-4e04-ae3d-1c53305d5fe6']
    c_res = run_command(cmd)
    logger.error("run_command(cmd) res:%s", c_res)

    # examples on how to read existing metrics from data_store
    m_res = data_store['metrics']['nsx.host.host-metrics']
    # examples on how to read effective profile from data_store
    profile = data_store['profile']
    logger.error("data_store['metrics']['nsx.host.host-metrics']:%s, profile:%s", m_res,
profile)

    # examples on how to read host_id from data_store
    host_id = data_store['host_id']
    logger.info('host_id:{}'.format(host_id))

    if profile['EXPORT_DATA']:
        report_test_data('running')
        logger.info("dynamic plugin - exported data to wavefront")
    else:
        logger.info("dynamic plugin - not export data ")

    # examples on how to use substitute interfaces for sys.
    logger.info("sys_path:{}".format(sys_path))
    logger.info("sys_getsizeof(1):{}".format(sys_getsizeof(1)))
    logger.info("sys_getrefcount(cmd):{}".format(sys_getrefcount(cmd)))
    logger.info("sys_maxsize:{}".format(sys_maxsize))
    logger.info("sys_getrecursionlimit():{}".format(sys_getrecursionlimit()))

    # examples on how to use substitute interfaces for datetime.
    today = datetime_date_today()
    logger.info("datetime today:{}".format(today))
    logger.info("datetime_date_strftime now:
{}".format(datetime_date_strftime(datetime.now(), '%H:%M')))
    logger.info('date.ctime(today):{}'.format(date.ctime(today)))


run()
```

Working with Events and Alarms

NSX provides alarms to call your attention to events that can potentially affect performance and system operation. Alarms provide detailed event information such as which component is affected, the type of event, and then recommends a corrective action.

For example, one of the NSX Edge nodes can be experiencing unusually high CPU usage or low available disk space.

Note Alarms are system events with a severity level greater than `LOW`.

If an alarm (for example, Certificate About to Expire) is raised and later a higher-severity alarm (for example, Certificate Expired) is raised about the same issue, the lower-severity alarm is not automatically resolved. You must take the recommended action to resolve the alarm.

Alarm information displays in several locations within the NSX Manager interface. For a complete list of events, see NSX Event Catalog.

View Alarm Information

Alarms information is displayed in several locations within the NSX Manager interface. Alarm and event information is also included with other notifications in the Notifications drop-down menu in the title bar.

An alarm can be in one of the following states:

State	Description
Open	Alarm is in an active, unacknowledged state.
Acknowledged	Alarm has been acknowledged by a user. The alarm remains open but no longer appears in the NSX Manager notifications.
Suppressed	Status reporting for this alarm has been disabled by the user for a user-specified duration.
Resolved	Alarm has been resolved, whether by the system or through user action. The alarm will continue to appear in the alarm table in the Resolved state for up to eight days, after which it automatically deletes. (The system may delete resolved alarms earlier to accommodate resource needs.) **Note** If a user changes an alarm state to Resolved but the condition that triggered the alarm is not resolved, a new alarm instance will be instantiated. Also, an event may be resolved for several minutes before the reported state updates in the interface.

Note The following steps show how to view alarms from the Home page. However, you can also view alarms from other pages, such as the Tier-0, Tier-1, and Load Balancing pages, among others. See the Alarms columns in the tables on these pages.

Procedure

1 Navigate to the Home page and click **Alarms**.

Note A red exclamation mark (**!**) next to the **Alarms** panel label indicates at least one open alarm with a severity of Critical. .

The Alarms panel appears, displaying along the top graphic dashboards such as Active Alarms, Top Features with the Most Alarms, and Top Events by Occurrence. Below the dashboards is a sortable, filterable list of the current alarms. The table details the following information about each active alarm:

- Feature affected
- Event Type
- Node
- Entity
- Severity (Critical, High, Medium)
- Last Reported Time
- Alarm State (Open, Suppressed, Resolved, Acknowledged)

Each row in the Alarms table can be expanded to show more details.

2 Filter the results displayed in the dashboards by clicking the funnel icon in the upper-right corner of the dashboards.

You can filter by the last 24 hours, last 48 hours, or custom time range, or all open alarms.

3 Filter the results displayed in the table by clicking the filter text box above the table.

You are prompted to specify a filter: Alarm State, Description, Entity Name, Entity Type, Event Type, Node, and so on.

What to do next

After viewing an alarm, you can decide on how to respond. See Managing Alarm States.

View Alarm Definitions

Detailed alarm definitions are provided on a separate panel in the Alarms tab. You can open the panel directly or arrive there by clicking the value in the Event Type column in the Alarms table.

Alarms details are displayed in several locations in the NSX Manager. See View Alarm Information.

Procedure

1 From the Alarms tab.

   a Navigate to the Home page and click **Alarms**.

   The Alarms panel has two modes, as shown at the top of the panel: **Alarms** and **Alarm Definitions**.

   b Click **Alarm Definitions**.

   The Alarms tab redisplays to show the table of Alarm definitions.

2 From the Tier-0 Gateways page.

   a Go to **Networking > Connectivity > Tier-0 Gateways**.

   The Open Alarms column of the gateway table displays the number of open alarms.

   b Click the number in the Open Alarms column.

   A dialog opens to display the open alarms in table format.

   c Click the value in the Event Type column.

   This action moves you to the **Home > Alarms > Alarm Definitions** panel described above.

3 From the Load Balancing page.

   a Go to **Networking > Network Services > Load Balancing**.

   The Open Alarms column of the gateway table displays the number of open alarms.

   b Click the number in the Open Alarms column.

   A dialog opens to display the open alarms in table format.

   c Click the value in the Event Type column.

   This action moves you to the **Home > Alarms > Alarm Definitions** panel described above.

4 After you access the **Alarm Definitions**, expand any definition to view details and user-definable settings.

   Alarm definition details include:

Column	Description
Feature	Displays the component where the alarm is originating, for example: Transport Node.
Event Type	Displays the specific type of error, for example: CPU Usage High.
Severity	Displays the level of alarm: Critical, High, or Medium.
Enabled	Displays whether detection of the Alarm is enabled.
Create Alarms	Displays whether to report the alarm in the interface or API.
Create SNMP Traps	Displays whether the system emits an SNMP trap when the alarm is detected or resolved.

   The panel also displays the following:

Item	Description
Description	Describes the condition that triggers the alarm.
Recommended Action	Describes steps you can take to correct the condition.
SNMP OID for Event true	Displays the SNMP Object Identifier for the Event when status is true.
SNMP OID for Event false	Displays the SNMP Object Identifier for the Event when status is false.
Threshold	User-configured threshold for triggering the alarm.
Sensitivity (%)	User-configured sensitivity for triggering the alarm.

What to do next

Some of fields in an alarm definition can be modified. See Configuring Alarm Definition Settings.

Configuring Alarm Definition Settings

Several settings in an alarm definition can be customized. From the Alarm Definitions page, you can enable or disable an alarm, configure if an event (when true) creates an alarm, create an SNMP trap, set alarm threshold, and set alarm sensitivity. From the Alarm Definitions page, you can enable or disable detection of an alarm, whether an alarm is reported in the API/user interface, and whether a SNMP trap is emitted when an alarm is detected or resolved.

You can configure the following alarm definition settings:

Setting	Control Type	Description
Enabled	Toggle	Enables or disables detection of the alarm.
Create Alarms	Toggle	Enables or disables whether the alarm is reported in the API/UI.
Create SNMP Traps	Toggle	Enables or disables whether an SNMP trap is emitted when an alarm is detected or resolved.
Threshold	Numerical value	Configures the threshold for triggering the event. This value determines if a single sample is true and triggers an event. ■ For CPU, disk, and memory alarms, threshold is the percentage usage value to indicate an alarming condition. ■ For certificate or license expiration alarms, this is the number of days before expiration, including local password expiration.
Sensitivity (%)	Numerical value (percentage)	Configures the sensitivity for triggering the alarm. Sensitivity defines the conditions that trigger an alarm. (The sample size is internally defined and cannot be modified.) If the sample size is ten and sensitivity is set to 80%, then eight or more occurrences in the sample of ten raises the alarms. See the NSX-T Data Center REST API documentation.

Procedure

1 Navigate to the Home page and click **Alarms**.

The Alarms panel has two modes, as shown at the top of the panel: **Alarms** and **Alarm Definitions**.

2 Click **Alarm Definitions**.

The Alarms tab redisplays to show the Alarm Definitions panel.

3 Right-click three vertical dots icon in the leftmost column of an alarm, and select **Edit**.

The selected alarm definition expands to show the definition details, and puts the configurable settings into edit mode.

4 Modify the settings as desired.

5 Click **Save**.

What to do next

For details about alarm definitions, see View Alarm Definitions. For details about SNMP traps, see Simple Network Management Protocol (SNMP).

Managing Alarm States

In addition to correcting the underlying causes, you can manage alarms by modifying their states as reported in the Alarms list.

Triggered alarms can be in one of the following states: Open, Acknowledged, Suppressed, or Resolved.

Procedure

1 Navigate to the Home page and click **Alarms**.

The Alarms panel has two modes, as shown at the top of the panel: **Alarms** and **Alarm Definitions**.

2 Click the **Alarms** mode, if the panel is not already displayed.

The **Alarms** panel displays a list of all alarms, including Resolved alarms.

Note Resolved alarms continue to be listed for up to eight days after their resolution.

3 Locate the alarm in the table on the page, and select the checkbox in the leftmost column.

4 Click **Action** and select the desired action.

- If you change the state of an alarm to Acknowledged, this indicates that you are aware of, and have acknowledged, the alarm.
- If you move an alarm into a Suppressed state, you are prompted to specify the duration in hours. After the specified duration passes, the alarm state reverts to Open. However, if the system determines the condition has been corrected, the alarm state changes to Resolved.
- You can restore the state of an Acknowledged or Suppressed alarm to Open.
- You cannot change the state of a Resolved alarm.

The value in the Alarm State column updates accordingly.

Registering Notification Watchers

You can register a watcher that will receive notifications based on specific criteria.

To register a watcher, you can invoke the following API.

`POST /api/v1/notification-watcher`

After adding a watcher, you must register a notification_id (feature_name.notification_name) with a watcher_id and specify notifications that the watcher should receive. Note that without the registration, the watcher will not receive any notifications. Invoke the following APIs with the required request parameters to register notification_id and to specify notifications. For more information on NSX Notification APIs, see *NSX API Guide*.

- `PUT /api/v1/notification-watchers/<watcher-id>`
- `POST /api/v1/notification-watchers/<watcher-id>/notifications?action=add_uri_filters` with the following request parameters:
  - notification_id: A string identifying feature_name.notification_name to indicate a notification that watcher is interested in receiving for the URI identified by the feature_name.notification_name.
  - uri_filters: Optional list of URIs to filter notifications based on its policy path. When specifying uri_filters, you can also use * as a wildcard character instead of a specific value.

    For example, if the notification_id is group.change_name, the uri_filter pattern is /policy/api/v1/infra/domains/<domain>/groups/<group>. You can specify the pattern as /policy/api/v1/infra/domains/domain1/groups/group2 to get notifications specific to domain1 and group2. Alternatively, you can also specify the pattern /policy/api/v1/infra/domains/domain2/groups/* to get notifications for all groups in domain2 or specify it as /policy/api/v1/infra/domains/*/groups/* to get notifications for all groups in all domains.

The following table lists the feature names and their respective URIs.

Feature Name	Feature Description	Notification Name	Notification Description	URI
group	Notifications supported by NS Group feature.	change_notification	Group notification, <domain> identifies the domain name and <group> identifies group name.	/policy/api/v1/infra/domains/<domain>/groups/<group>
monitoring	Notifications supported by the monitoring feature.	alarm	Alarm notifications. <alarm-id> identifies an alarm instance. A notification is sent whenever an alarm instance is created or deleted and when the alarm instance is updated.	/api/v1/alarms/<alarm-id>

Feature Name	Feature Description	Notification Name	Notification Description	URI
		alarm_status_change_notification	Alarm notifications. <alarm-id> identifies an alarm instance. A notification is sent whenever an alarm instance is created and when the status property value of an alarm instance is updated.	/api/v1/alarms/<alarm-id>
notification	Notifications supported by notification framework.	watcher	Platform notification to convey updates to watcher configuration. <watcher-id> identifies the watcher.	/api/v1/notification-watchers/<watcher-id>
		watcher_notification	Platform notification to convey updates to notifications. <watcher-id> identifies the watcher.	/api/v1/notification-watchers/<watcher-id>/notifications
service_config	Notifications supported by Service Config feature.	change_notification	Service config notification. <domain> identifies the domain name, <policy> identifies the endpoint policy, and <rule> identifies the endpoint rule. This notification is generated when a service config used in endpoint rule is updated or when UPM Profile is updated.	/policy/api/v1/infra/domains/<domain>/endpoint-policies/<policy>/endpoint-rules/<rule>
service_insertion	Notifications supported by Service Insertion module. Currently Service Insertion module supports notifications for Service Profile, Service Instance Runtime, and Policy Groups.	instance_runtime_notification	Service Instance Runtime notification. <service-id> identifies the service, <service-instance-id> identifies the service instance. Notification will be sent for deployed and undeployed operations.	/api/v1/serviceinsertion/services/<service-id>/service-instances/<service-instance-id>/instance-runtimes
		profile_notification	Service Profile change notification. <service-reference> identifies the service name and <service-profile> identifies profile name. Notification will be sent for profile create, update, and delete.	/policy/api/v1/infra/service-references/<service-reference>/service-profiles/<service-profile>

Feature Name	Feature Description	Notification Name	Notification Description	URI
		profile_chain_mapping_notification	Service Profile Chain Mapping notification. <service-reference> identifies the service name and <service-profile> identifies profile name. The notification will be sent when a profile is added or removed as a part of a service chain.	/policy/api/v1/infra/service-references/<service-reference>/service-profiles/<service-profile>/service-chain-mappings
		profile_nsgroups_notification	Service Profile NSGroups notification. <service-reference> identifies the service name and <service-profile> identifies profile name. This notification gets triggered whenever an east-west rule containing nsgroups gets added or deleted with the particular profile.	/policy/api/v1/infra/service-references/<service-reference>/service-profiles/<service-profile>/group-associations
		instance_nsgroups_notification	Service Instance NSGroups notification. <service-id> identifies the service name and <service-instance-id> identifies service instance. This notification gets triggered whenever a north-south rule containing nsgroups gets added or deleted with the particular instance.	/api/v1/serviceinsertion/services/<service-id>/service-instances/<service-instance-id>/group-associations

Using Aria Operations for Logs or Splunk for System Monitoring

You can monitor your NSX environment using Aria Operations for Logs or Splunk.

You can find the NSX Splunk app at https://splunkbase.splunk.com/app/4241.

The Aria Operations for Logs content pack has the following alerts:

Alert Name	Description
SysCpuUsage	CPU usage is above 95% for more than 10 minutes.
SysMemUsage	Memory usage is above 95% for more than 10 minutes.
SysDiskUsage	Disk usage for one or more partitions is above 89% for more than 10 minutes.

Alert Name	Description
PasswordExpiry	Password for appliance user account is about to expire or expired.
CertificateExpiry	One or more CA signed certificate is expired.
ClusterNodeStatus	Local edge cluster node is down.
BackupFailure	NSX scheduled backup operation failed.
VipLeadership	NSX Management cluster VIP is down.
ApiRateLimit	Client API reached configured threshold.
CorfuQuorumLost	Two nodes went down in the cluster and lost corfu quorum.
DfwHeapMem	DFW heap memory exceeded configured threshold.
ProcessStatus	Critical process status changed.
ClusterFailoverStatus	SR high availability state changed or active/standby services failover.
DhcpPoolUsageOverloadedEvent	DHCP pool reached configured usage threshold.
FabricCryptoStatus	Edge crypto mux driver is down for failing Known_Answer_Tests (KAT).
VpnTunnelState	VPN tunnel is down.
BfdTunnelStatus	BFD Tunnel status changed.
RoutingBgpNeighborStatus	BGP neighbor status is down.
VpnL2SessionStatus	L2 VPN session is down.
VpnIkeSessionStatus	IKE session is down.
RoutingStatus	Routing(BGP/BFD) is down.
DnsForwarderStatus	DNS forwarder running status is DOWN.
TnConnDown_15min	Transport Node connection to a controller/Manager is down for at least 15 minutes.
TnConnDown_5min	Transport Node connection to controller/Manager is down for at least 5 minutes.
ServiceDown	One or more services are down.
IpNotAvailableInPool	There is no IP available in the Pool or reaches configured threshold.
LoadBalancerError	NSX Load Balancer Service status is ERROR.
LoadBalancerDown	NSXLoad Balancer Service status is DOWN.
LoadBalancerVsDown	VS status: all pool members are down.
LoadBalancerPoolDown	Pool status: all pool members are down.
ProcessCrash	Process or daemon crashes in the datapath or other LB process like dispatcher, etc..

Dashboards

Both the Splunk app and the Aria Operations for Logs content pack have the following dashboards.

Table 20-1. NSX - Infrastructure

Widget Name	Notes
NSX Manager: Communication Errors	These are all communication log errors on NSX Manager. They are grouped by hostname.
Transport Node - NSX Manager : Communication Errors	Communication errors between NSX transport nodes (vSphere hosts, KVM hosts and NSX Edges) and NSX Manager. It is recommended to analyze the hostnames with the highest returned values to look for potential issues.
Transport Node - Controller : Communication Errors	Communication errors between NSX transport nodes (vSphere hosts, KVM hosts and NSX Edges) and NSX controllers. It is recommended to analyze the hostnames with the highest returned values to look for potential issues.
Controller: Communication Errors	Communication errors among controllers in a cluster. It is recommended to analyze the hostnames with the highest returned values to look for potential issues.
Configuration Errors	This widget is based on known error patterns generated from various components that create the NSX infrastructure. It is recommended to analyze the hostnames with the highest returned values to look for potential issues.
Other Errors	Other errors from all NSX components. It is recommended to analyze the hostnames with the highest returned values to look for potential issues. Host agent failures, such as netcpa down or MPA down, may require immediate action.

Table 20-2. NSX - Audits

Widget Name	Notes
Logical Switch Audits	Logical switch messages excerpted from the NSX audit log capturing any create, update, or delete events. Note: Events may be duplicated as they are generated in multiple NSX loggers. Thus, this widget will provide a general count of logical switch audit events.
Logical Switch Audit Details	All logical switch audit events. This widget is based on known error patterns generated from various components that create the NSX logical switch infrastructure.
Logical Switch Port Audits	Logical switch port messages excerpted from the NSX audit log capturing any create, update, or delete events. Note: Events may be duplicated as they are generated in multiple NSX loggers. Thus, this widget will provide a general count of logical switch port audit events.
Logical Switch Port Audit Details	All logical switch port audit events. This widget is based on known error patterns generated from various components that create the NSX logical switch infrastructure.
Logical Router Audits	Logical router messages excerpted from the NSX audit log capturing any create, update, or delete events. Note: Events may be duplicated as they are generated in multiple NSX loggers. Thus, this widget will provide a general count of logical router audit events.
Logical Router Audit Details	All logical router audit events. This widget is based on known error patterns generated from various components that create the NSX logical routing infrastructure.
Logical Router Port Audits	Logical router port messages excerpted from the NSX audit log capturing any create, update, or delete events. Note: Events may be duplicated as they are generated in multiple NSX loggers. Thus, this widget will provide a general count of logical router port audit events.
Logical Router Port Audit Details	All logical switch port audit events. This widget is based on known error patterns generated from various components that create the NSX logical router infrastructure.

Table 20-2. NSX - Audits (continued)

Widget Name	Notes
Firewall Audits	Firewall messages excerpted from the NSX audit log capturing any add section, update section with new rules, or delete section events. Note: Events may be duplicated as they are generated in multiple NSX loggers. Thus, this widget will provide a general count of firewall audit events.
Logical Firewall Audit Details	All firewall audit events. This widget is based on known error patterns generated from the NSX firewall.

Table 20-3.

MSX - Logical Switch	Notes
Logical Switch Created	Logical switch messages excerpted from the NSX audit log capturing any create events. Note: Events may be duplicated as they are generated in multiple NSX loggers. Thus, this widget will provide a general count of logical switch audit events.
Logical Switch Updates	Logical switch messages excerpted from the NSX audit log capturing any update events. Note: Events may be duplicated as they are generated in multiple NSX loggers. Thus, this widget will provide a general count of logical switch audit events.
Logical Switch Deleted	Logical switch messages excerpted from the NSX audit log capturing any delete events. Note: Events may be duplicated as they are generated in multiple NSX loggers. Thus, this widget will provide a general count of logical switch audit events.
Logical Switch Audit Details	All logical switch audit events.
Logical Switch - Manager Errors	Logical switch log errors reported to NSX Manager. Errors are grouped by hostname.
Logical Switch - Controller Errors	Logical switch log errors reported from the view of the NSX controllers. Errors are grouped by hostname.
Logical Switch - Transport Node Errors	Logical switch log errors reported for all NSX transport nodes - vSphere hosts, KVM hosts, and Edge Services Gateways. Errors are grouped by node.

Table 20-4. NSX - Logical Router

Widget Name	Notes
Logical Router Create Audit Events	Logical router messages excerpted from the NSX audit log capturing any create events. Note: Events may be duplicated as they are generated in multiple NSX loggers. Thus, this widget will provide a general count of logical router audit events.
Logical Router Update Audit Events	Logical router messages excerpted from the NSX audit log capturing any update events. Note: Events may be duplicated as they are generated in multiple NSX loggers. Thus, this widget will provide a general count of logical router audit events.
Logical Router Delete Audit Events	Logical router messages excerpted from the NSX audit log capturing any delete events. Note: Events may be duplicated as they are generated in multiple NSX loggers. Thus, this widget will provide a general count of logical router audit events.
Logical Router Audit Details	All logical router audit events.
Logical Router - Manager Errors	Logical router log errors reported to NSX Manager. Errors are grouped by hostname.

Table 20-4. NSX - Logical Router (continued)

Widget Name	Notes
Logical Router - Controller Errors	Logical router log errors detected by NSX controllers. Errors are grouped by hostname.
Logical Router - Transport Node Errors	Logical router log errors for NSX transport nodes - vSphere hosts, KVM hosts, and Edge Services Gateways. Errors are grouped by node.

Table 20-5. NSX - Distributed Firewall Overview

Widget Name	Notes
Section Create Events	All firewall section create audit events. Note: Events may be duplicated as they are generated in multiple NSX loggers. Thus, this widget will provide a general count of firewall changes.
Section Update Events	All firewall section update audit events. Note - Any create, update, delete activity on a rule raising section update event. Note: Events may be duplicated as they are generated in multiple NSX loggers. Thus, this widget will provide a general count of firewall changes.
Section Delete Events	All firewall section delete audit events. Note: Events may be duplicated as they are generated in multiple NSX loggers. Thus, this widget will provide a general count of firewall changes.
Section Audit Details	Shows all Firewall audit events (who changed what). Note - Any create, update, delete activity on a rule raising section update event.
Firewall - Manager Errors	Firewall log errors reported to NSX Manager. Errors are grouped by hostname.
Firewall - Controller Errors	Firewall log errors reported to NSX controllers. Errors are grouped by hostname.
Firewall - Transport Node Errors	Firewall log errors on NSX transport nodes - vSphere hosts and KVM hosts. Errors are grouped by node.

Table 20-6. NSX - Distributed Firewall Traffic

Widget Name	Notes
Top Firewall Sources	Top source IP addresses from all firewall rules that are logging data.
Top Firewall Destinations	Top destination IP addresses from all firewall rules that are logging data.
Application Ports Permitted	Measures all in / out connections permitted in the NSX-T environment by destination port. The data is the summation of the specified time range.
Application Ports Denied	All traffic defined by a firewall rule. Data is grouped by application (or destination) port number. This widget displays only data associated with a port. Traffic types, such as ICMP, without an associated port are not displayed.
Top Firewall Sources by bytes - client to server	All firewall traffic, in bytes by source IP address, from the client to a server. Data is only displayed if the firewall is logging its data. The data is the summation of the specified time range.
Top Firewall Destinations by bytes - client to server	All firewall traffic, in bytes by destination IP address, from the client to a server. Data is only displayed if the firewall is logging its data. The data is the summation of the specified time range.

Table 20-6. NSX - Distributed Firewall Traffic (continued)

Widget Name	Notes
Top Firewall Sources by bytes - server to client	All firewall traffic, in bytes by source IP address, from the server to a client. Data is only displayed if the firewall is logging its data. The data is the summation of the specified time range.
Top Firewall Destinations by bytes - server to client	All firewall traffic, in bytes by destination IP address, from the server to a client. Data is only displayed if the firewall is logging its data. The data is the summation of the specified time range.

Table 20-7. NSX - DHCP

Widget Name	Notes
DHCP Create Audit Events	All DHCP create audit events, including new DHCP profiles, static bindings, or IP pools. Note: Events may be duplicated as they are generated in multiple NSX loggers. Thus, this widget will provide a general count of firewall changes.
DHCP Update Audit Events	All DHCP update audit events, including updated DHCP profiles, static bindings, and IP pools. Note: Events may be duplicated as they are generated in multiple NSX loggers. Thus, this widget will provide a general count of firewall changes.
DHCP Delete Audit Events	All DHCP delete audit events, including deleted DHCP profiles, static bindings, or IP pools. Note: Events may be duplicated as they are generated in multiple NSX loggers. Thus, this widget will provide a general count of firewall changes.
DHCP Audit Details	All DHCP audit events.
DHCP - Manager Errors	DHCP log errors reported by NSX Manager.
DHCP - Controller Errors	DHCP log errors reported by NSX controllers.
DHCP - Transport Node Errors	DHCP log errors reported for NSX transport nodes - vSphere hosts, KVM hosts, and Edge Services Gateways.

Table 20-8. NSX - Backup

Widget Name	Notes
Backup Configuration Updates	Total number of times backup configuration was updated.
Backup Failures	Count of all backup failures over time grouped by error code.
Successful Cluster Backups	Total number of cluster and node backups completed successfully.
Failed Cluster Backups	Total number of cluster and node backups failed.
Successful Inventory Backups	Total number of inventory backups completed successfully.
Failed Inventory Backups	Total number of inventory backups failed.

Table 20-9. NSX - IPAM

Widget Name	Notes
IPAM Create Events	IPAM messages excerpted from the NSX audit log capturing any create events. Note: Events may be duplicated as they are generated in multiple NSX loggers. Thus, this widget will provide a general count of IPAM audit events.
IPAM Update Events	IPAM messages excerpted from the NSX audit log capturing any update events. Note: Events may be duplicated as they are generated in multiple NSX loggers. Thus, this widget will provide a general count of IPAM audit events.
IPAM Delete Events	IPAM messages excerpted from the NSX audit log capturing any delete events. Note: Events may be duplicated as they are generated in multiple NSX loggers. Thus, this widget will provide a general count of IPAM audit events.
IPAM Audit Details	All IPAM audit events.
IPAM - Manager Errors	IPAM log errors reported by NSX Manager.

Table 20-10. NSX - Unified Security Flow Logs

Widget Name	Notes
Top Security Vertical Source	Top source IP addresses from all security verticals that are logging data.
Top Security Vertical Destination	Top destination IP addresses from all security verticals that are logging data.
Application Ports Permitted	Measures all in / out connections permitted in the NSX-T environment by destination port. The data is the summation of the specified time range.
Application Ports Denied	All security flows denied by a security vertical rule. Data is grouped by application (or destination) port number. This widget displays only data associated with a port. Traffic types, such as ICMP, without an associated port are not displayed.
Top Security Vertical Sources by bytes - client to server	All security flows, in bytes by source IP address, from the client to a server. Data is only displayed if the vertical is logging its data. The data is the summation of the specified time range.
Top Security Vertical Destinations by bytes - client to server	All security flows, in bytes by destination IP address, from the client to a server. Data is only displayed if the vertical is logging its data. The data is the summation of the specified time range.
Top Security Vertical Sources by bytes - server to client	All security flows, in bytes by source IP address, from the server to a client. Data is only displayed if the vertical is logging its data. The data is the summation of the specified time range.
Top Security Vertical Destinations by bytes - server to client	All security flows, in bytes by destination IP address, from the server to a client. Data is only displayed if the vertical is logging its data. The data is the summation of the specified time range.

Using Aria Operations Manager for System Monitoring

You can monitor your NSX environment using Aria Operations Manager.

Table 20-11. Alerts in the Management Pack for NSX

Alert	Description	Recommendation
NSX Management service has failed	Triggered when the management service on the NSX host is not running.	Log in to the NSX Manager and restart the failed management service.
Logical Switch's admin state is not UP	Triggered when the admin state is disabled on the logical switch.	Log in to NSX and enable the admin state if it is intended so.
Edge Node Controller/ Manager Connectivity is not UP	Triggered when the edge node connectivity status is down in NSX.	Check the Edge node connection status with Controller Cluster and Manager Cluster and fix the broken connection.
Edge Host node is in Failed/Error state	Triggered when the host node in NSX is in error or failed state due to one of the following reasons: ■ Edge configuration error ■ Installation failure ■ Uninstallation failure ■ Upgrade failure ■ Virtual Machine deployment failure ■ Virtual Machine power off failure ■ Virtual Machine power on failure ■ Virtual Machine undeployment failure	Edge host node is in failed/ error state, check the host node state and fix the issue.
BFD service is disabled	Triggered when the BFD service is not enabled on the logical router.	BFD Service for a TIER0 router is not enabled even though neighbors are configured. Enable the BFD service if required.
NAT rule not configured	Triggered when the NAT rule on the logical router is not configured.	Log in to the NSX Manager and add the NAT rules for the Logical Router.
Static Route not configured	Triggered when the static route on the logical router is not configured.	Log in to the NSX Manager and add the static routes for the Logical Router if required.
Route Advertisement service is disabled	Triggered when the route advertisement service is not enabled on the logical router.	Route Advertisement service for a TIER1 router is not enabled even though route advertisements are configured, log in to NSX Manager and enable the service.

Table 20-11. Alerts in the Management Pack for NSX (continued)

Alert	Description	Recommendation
Route Redistribution service is disabled	Triggered when the route redistribution service is not enabled on the logical router.	Route Redistribution service for a TIER0 router is not enabled even though route redistribution rules are configured, log in to NSX Manager and enable the service.
ECMP service is disabled for Logical Router	Triggered when the ECMP service is not enabled on the logical router.	BGP ECMP service for a TIER0 router is not enabled even though neighbors are configured, log in to NSX Manager and enable the service.
Controller Node Connectivity is broken	Triggered when the controller node connection status is down in NSX	Log in to NSX Manager and check the connectivity of the controller node with Management Node and Controller cluster and resolve the disconnected state.
Less than 3 controller nodes are deployed	Triggered when the NSX server has less than three controller nodes.	Deploy at least 3 controller nodes in the cluster.
Controller Cluster Status is not stable	Triggered when all the controller nodes are down in NSX.	Check the status of controller cluster.
Management Status is not stable	Triggered when the status of any node on the management cluster is down.	Check the status of management cluster.
File System usage is more than 85 percent	Triggered when the guest file systems usage of the Controller Virtual Machine is more than 85 percent.	File system usage is more than 85, check and clean the File System to make more space.
File System usage is more than 75 percent	Triggered when the guest file systems usage of the Controller Virtual Machine is more than 75 percent.	File system usage is more than 75, check and clean the File System to make more space.
File System usage is higher than 70 percent	Triggered when the guest file systems usage of the Controller Virtual Machine is more than 70 percent.	File system usage is more than 70, check and clean the File System to make more space.
Edge Cluster Status is down	Triggered when edge cluster status is down.	Check the edge cluster status and if required follow standard troubleshooting steps recommended by NSX documentation and VMware documentation.

Table 20-11. Alerts in the Management Pack for NSX (continued)

Alert	Description	Recommendation
Logical Switch State has failed	Triggered when the state of logical switch has failed.	Check the logical switch state and if necessary follow standard troubleshooting steps recommended by NSX documentation and VMware documentation.
Load Balancer Service operational status down	Triggered when the operational status of load balancer service is down.	Check the operational status of load balancer service and if necessary follow standard troubleshooting steps recommended by NSXdocumentation and VMware documentation.
Load balancer service operational status error	Triggered when the operational status of load balancer service contains error.	Check the operational status of load balancer service and if necessary follow standard troubleshooting steps recommended by NSX documentation and VMware documentation.
Load Balancer virtual server operational state down	Triggered when the operational state of load balancer virtual server is down.	Check the operational state of load balancer virtual server and if necessary follow standard troubleshooting steps recommended by NSX documentation and VMware documentation.
Load Balancer virtual server operational state detached	Triggered when the operational state of load balancer virtual server is detached.	Check the operational state of load balancer virtual server and if necessary follow standard troubleshooting steps recommended by NSX documentation and VMware documentation.
Edge node configuration state has failed	Triggered when the configuration state of edge node has failed.	Check the configuration state of the edge node and if necessary follow standard troubleshooting steps recommended by NSX documentation and VMware documentation.

Table 20-11. Alerts in the Management Pack for NSX (continued)

Alert	Description	Recommendation
Management service monitor runtime state has failed	Triggered when the monitor runtime state of the management service stops running.	Log in to the NSX Manager VA and restart the failed management service.
Management cluster's management status is not stable	Triggered when the management status of a management cluster is not stable.	Check the status of management cluster.
Less than 3 manager nodes are deployed	Triggered when the NSX server has less than three manager nodes deployed.	Deploy at least 3 manager nodes in the cluster.
Manager node connectivity is broken	Triggered when the manager connection status of manager node is down.	Log in to NSX Manager and check the manager connectivity of manager node and follow standard troubleshooting steps recommended by NSXdocumentation and VMware documentation.
File System usage of manager node is more than 85 percent	Triggered when the guest file systems usage of the manager node is more than 85 percent.	File system usage is more than 85, check and clean the File System to make more space.
File System usage of manager node is more than 75 percent	Triggered when the guest file systems usage of the manager node is more than 75 percent.	File system usage is more than 75, check and clean the File System to make more space.
File System usage of manager node is more than 70 percent	Triggered when the guest file systems usage of the manager node is more than 70 percent.	File system usage is more than 70, check and clean the File System to make more space.

Using Aria Operations for Networks for System Monitoring

You can monitor your NSX environment using Aria Operations for Networks (formerly known as vRealize Network Insight Cloud).

Table 20-12. Aria Operations for Networks Computed NSX Events

OID	Event Name	Default Severity	UI Name	Description
1.3.6.1.4.1.6876.100.1.0.80205	NSXTNoUplinkConnectivityEvent	Warning	NSX Tier-1 logical router disconnect event	NSX Tier-1 logical router is disconnected from Tier-0 router. Networks under this router are not reachable from outside and vice versa.
1.3.6.1.4.1.6876.100.1.0.80206	NSXTRoutingAdvertisementEvent	Warning	Routing advertisement disabled	Routing advertisement is disabled for NSX Tier-1 logical router. Networks under this router are not reachable from outside.
1.3.6.1.4.1.6876.100.1.0.80207	NSXTManagerConnectivityDownEvent	Critical	NSX Edge Node has no manager connectivity	NSX Edge Node has lost manager connectivity.
1.3.6.1.4.1.6876.100.1.0.80208	NSXTControllerConnectivityDegradedEvent	Warning	Controller connectivity degraded for NSX Edge Node	NSX Edge Node is not able to communicate with one or more controllers.
1.3.6.1.4.1.6876.100.1.0.80209	NSXTControllerConnectivityDownEvent	Critical	NSX Edge Node has no controller connectivity	NSX Edge Node is not able to communicate with any of the controllers.
1.3.6.1.4.1.6876.100.1.0.80210	NSXTMtuMismatchEvent	Warning	MTU mismatch between NSX Tier-0 and uplink switch/ router	The MTU configured on interfaces of Tier-0 logical router does not match with the interfaces of uplink switch/router from same L2 network. This can impact the network performance.
1.3.6.1.4.1.6876.100.1.0.80211	NSXTExcludedVmFlowEvent	Info	One or More VMs excluded from NSX DFW Firewall.	One or more VMs are not protected by NSX DFW firewall. Aria Operations for Networks will not receive IPFIX flows for these VMs.

Table 20-12. Aria Operations for Networks Computed NSX Events (continued)

OID	Event Name	Default Severity	UI Name	Description
1.3.6.1.4.1.6876.100.1.0.80212	NSXTDoubleVlanTaggingEvent	Warning	Uplink Vlan misconfiguration	Communication is disrupted because VLAN on uplink port of Tier 0 router is different than VLAN on the external gateway.
1.3.6.1.4.1.6876.100.1.0.80213	NSXTNoTzAttachedOnTnEvent	Warning	No transport zone is attached to the transport node.	No transport zone attached to the transport node. VMs might lose connectivity because of this.
1.3.6.1.4.1.6876.100.1.0.80214	NSXTVtepDeleteEvent	Warning	No VTEP available on the transport node.	All vteps are deleted from the transport node. VMs might lose connectivity because of this.
1.3.6.1.4.1.6876.100.1.0.80225	NSXTControllerNodeToControlClusterConnectivityEvent	Critical	NSX controller node has no control cluster connectivity	NSX controller node has lost control cluster connectivity.
1.3.6.1.4.1.6876.100.1.0.80226	NSXTControllerNodeToMgmtPlaneConnectivityEvent	Critical	NSX controller node has no management plane connectivity	NSX controller node has lost management plane connectivity.
1.3.6.1.4.1.6876.100.1.0.80227	NSXTMPNodeToMgmtClusterConnectivityEvent	Critical	NSX management node has no management cluster connectivity	NSX management node has lost management cluster connectivity.
1.3.6.1.4.1.6876.100.1.0.80246	NSXTHostNodeMgmtConnectivityStatusDownEvent	Warning	NSX Host Node has no manager connectivity	Desynchronization between NSX Manager's State of connectivity with Host Transport Nodes
1.3.6.1.4.1.6876.100.1.0.80247	NSXTEdgeNodeCtlrConnectivityStatusUnknownEvent	Critical	Controller connectivity for NSX Edge Node is Unknown.	NSX Edge Node Controller connectivity is Unknown.
1.3.6.1.4.1.6876.100.1.0.80248	NSXTHostNodeCtlrConnectivityStatusDownEvent	Warning	NSX Host Node has no controller connectivity	NSX Host Node is not able to communicate with any of the controllers.
1.3.6.1.4.1.6876.100.1.0.80249	NSXTHostNodeCtlrConnectivityStatusDegradedEvent	Warning	Controller connectivity degraded for NSX Host Node	NSX Host Node is not able to to communicate with one or more controllers.

Table 20-12. Aria Operations for Networks Computed NSX Events (continued)

OID	Event Name	Default Severity	UI Name	Description
1.3.6.1.4.1.6876.100.1.0.80250	NSXTHostNodeCtlrConnectivityStatusUnknownEvent	Warning	Controller connectivity for NSX Host Node is Unknown.	NSX Host Node Controller connectivity is Unknown.
1.3.6.1.4.1.6876.100.1.0.80228	NSXTHostNodePnicStatusDownEvent	Warning	NSX Host Transport Node Pnic Status is 'Down'.	NSX Host Transport Node Pnic Status is 'Down'.
1.3.6.1.4.1.6876.100.1.0.80229	NSXTHostNodePnicStatusDegradedEvent	Warning	NSX Host Transport Node Pnic Status is 'Degraded'	NSX Host Transport Node Pnic Status is 'Degraded'.
1.3.6.1.4.1.6876.100.1.0.80230	NSXTHostNodePnicStatusUnknownEvent	Warning	NSX Host Transport Node Pnic Status is 'Unknown'.	NSX Host Transport Node Pnic Status is 'Unknown'.
1.3.6.1.4.1.6876.100.1.0.80237	NSXTEdgeNodePnicStatusDownEvent	Critical	NSX Edge Transport Node Pnic Status is 'Down'.	NSX Edge Transport Node Pnic Status is 'Down'.
1.3.6.1.4.1.6876.100.1.0.80238	NSXTEdgeNodePnicStatusDegradedEvent	Critical	NSX Edge Transport Node Pnic Status is 'Degraded'.	NSX Edge Transport Node Pnic Status is 'Degraded'.
1.3.6.1.4.1.6876.100.1.0.80239	NSXTEdgeNodePnicStatusUnknownEvent	Critical	NSX Edge Transport Node Pnic Status is 'Unknown'.	NSX Edge Transport Node Pnic Status is 'Unknown'.
1.3.6.1.4.1.6876.100.1.0.80231	NSXTHostNodeTunnelStatusDownEvent	Warning	NSX Host Transport Node Tunnel Status is 'Down'.	NSX Host Transport Node Tunnel Status is 'Down'.
1.3.6.1.4.1.6876.100.1.0.80232	NSXTHostNodeTunnelStatusDegradedEvent	Warning	NSX Host Transport Node Tunnel Status is 'Degraded'.	NSX Host Transport Node Tunnel Status is 'Degraded'.
1.3.6.1.4.1.6876.100.1.0.80233	NSXTHostNodeTunnelStatusUnknownEvent	Warning	NSX Host Transport Node Tunnel Status is 'Unknown'.	NSX Host Transport Node Tunnel Status is 'Unknown'.
1.3.6.1.4.1.6876.100.1.0.80240	NSXTEdgeNodeTunnelStatusDownEvent	Critical	NSX Edge Transport Node Tunnel Status is 'Down'.	NSX Edge Transport Node Tunnel Status is 'Down'.
1.3.6.1.4.1.6876.100.1.0.80241	NSXTEdgeNodeTunnelStatusDegradeEvent	Critical	NSX Edge Transport Node Tunnel Status is 'Degraded'.	NSX Edge Transport Node Tunnel Status is 'Degraded'.
1.3.6.1.4.1.6876.100.1.0.80242	NSXTEdgeNodeTunnelStatusUnknownEvent	Critical	NSX Edge Transport Node Tunnel Status is 'Unknown'.	NSX Edge Transport Node Tunnel Status is 'Unknown'.
1.3.6.1.4.1.6876.100.1.0.80234	NSXTHostNodeStatusDownEvent	Warning	NSX Host Transport Node Status is 'Down'.	NSX Host Transport Node Status is 'Down'.

Table 20-12. Aria Operations for Networks Computed NSX Events (continued)

OID	Event Name	Default Severity	UI Name	Description
1.3.6.1.4.1.6876.100.1.0.80235	NSXTHostNodeStatusDegradedEvent	Warning	NSX Host Transport Node Status is 'Degraded'.	NSX Host Transport Node Status is 'Degraded'.
1.3.6.1.4.1.6876.100.1.0.80236	NSXTHostNodeStatusUnknownEvent	Warning	NSX Host Transport Node Status is 'Unknown'.	NSX Host Transport Node Status is 'Unknown'.
1.3.6.1.4.1.6876.100.1.0.80243	NSXTEdgeNodeStatusDownEvent	Critical	NSX Edge Transport Node Status is 'Down'.	NSX Edge Transport Node Status is 'Down'.
1.3.6.1.4.1.6876.100.1.0.80244	NSXTEdgeNodeStatusDegradedEvent	Critical	NSX Edge Transport Node Status is 'Degraded'.	NSX Edge Transport Node Status is 'Degraded'.
1.3.6.1.4.1.6876.100.1.0.80245	NSXTEdgeNodeStatusUnknownEvent	Critical	NSX Edge Transport Node Status is 'Unknown'.	NSX Edge Transport Node Status is 'Unknown'.
1.3.6.1.4.1.6876.100.1.0.80252	NSXTLogicalSwitchAdminStatusDownEvent	Warning	NSX Logical Switch Admin Status is 'Down'	NSX Logical Switch Admin Status is 'Down'
1.3.6.1.4.1.6876.100.1.0.80253	NSXTLogicalPortOperationalStatusDownEvent	Critical	NSX Logical Port Operational Status is 'Down'	NSX Logical Port Operational Status is 'Down'. This could cause a communication failure between two virtual interfaces (VIFs) that are connected to the same logical switch, for example, you cannot ping one VM from another.
1.3.6.1.4.1.6876.100.1.0.80254	NSXTLogicalPortOperationalStatusUnknownEvent	Warning	NSX Logical Port Operational Status is 'Unknown'	NSX Logical Port Operational Status is 'Unknown'. This could cause a communication failure between two virtual interfaces (VIFs) that are connected to the same logical switch, for example, you cannot ping one VM from another.

Table 20-12. Aria Operations for Networks Computed NSX Events (continued)

OID	Event Name	Default Severity	UI Name	Description
1.3.6.1.4.1.6876.100.1.0.80255	NSXTComputeManagerConnectionStatusNotUpEvent	Warning	NSX Compute Manager Connection Status in not up	NSX Compute Manager Connection status is not up
1.3.6.1.4.1.6876.100.1.0.80256	NSXTClusterBackUpDisabledEvent	Warning	NSX Manager backup is not scheduled.	NSX Manager backup is not scheduled
1.3.6.1.4.1.6876.100.1.0.80257	NSXTDFWFirewallDisabledEvent	Critical	NSX DFW Firewall is disabled.	Distributed Firewall is disabled in the NSX Manager
1.3.6.1.4.1.6876.100.1.0.80258	NSXTLogicalPortReceivedPacketDropEvent	Warning	NSX Logical Port Received Packets are getting dropped.	Received packets are getting dropped on the NSX Logical Port and associated entities might get affected
1.3.6.1.4.1.6876.100.1.0.80259	NSXTLogicalPortTransmittedPacketDropEvent	Warning	NSX Logical Port Transmitted Packets are getting dropped.	Transmitted packets are getting dropped on the NSX Logical Port and associated entities might get affected
1.3.6.1.4.1.6876.100.1.0.80260	NSXTLogicalSwitchReceivedPacketDropEvent	Warning	NSX Logical Switch Received Packets are getting dropped	Received packets are getting dropped on the NSX Logical Switch and associated entities might get affected
1.3.6.1.4.1.6876.100.1.0.80261	NSXTLogicalSwitchTransmittedPacketDropEvent	Warning	NSX Logical Switch Transmitted Packets are getting dropped	Transmitted packets are getting dropped on the NSX Logical Switch and associated entities might get affected
1.3.6.1.4.1.6876.100.1.0.80262	NSXTRxPacketDropOnMPNicEvent	Warning	Received packets are dropping on NSX Management Node's network interface	Received packets are getting dropped on NSX Management Node's network interface. This may impact the network traffic related to NSX management cluster.
1.3.6.1.4.1.6876.100.1.0.80263	NSXTRxPacketDropOnEdgeTnNicEvent	Critical	Received packets are dropping on NSX Edge Node's network interface	Received packets are getting dropped on NSX Edge Node's network interface. This may impact the network traffic of edge cluster.

Table 20-12. Aria Operations for Networks Computed NSX Events (continued)

OID	Event Name	Default Severity	UI Name	Description
1.3.6.1.4.1.6876.100.1.0.80264	NSXTRxPacketDropOnHostTnNicEvent	Warning	Received packets are dropping on NSX Host Node's network interface	Received packets are getting dropped on NSX Host Node's network interface. This may impact the network traffic on ESXi Host.
1.3.6.1.4.1.6876.100.1.0.80265	NSXTTxPacketDropOnMPNicEvent	Warning	Transmitted packets are dropping on NSX Management Node's network interface	Transmitted packets are getting dropped on NSX Management Node's network interface. This may impact the network traffic related to NSX management cluster.
1.3.6.1.4.1.6876.100.1.0.80266	NSXTTxPacketDropOnEdgeTnNicEvent	Critical	Transmitted packets are dropping on NSX Edge Node's network interface	Transmitted packets are getting dropped on NSX Edge Node's network interface. This may impact the network traffic of edge cluster.
1.3.6.1.4.1.6876.100.1.0.80267	NSXTTxPacketDropOnHostTnNicEvent	Warning	Transmitted packets are dropping on NSX Host Node's network interface	Transmitted packets are getting dropped on NSX Host Node's network interface. This may impact the network traffic on ESXi Host.
1.3.6.1.4.1.6876.100.1.0.80402	NSXTMPNodeServiceCmInventoryStatusEvent	Warning	CM Inventory Serivce has stopped running	CM Inventory Service status has turned to stopped.
1.3.6.1.4.1.6876.100.1.0.80403	NSXTMPNodeServiceControllerStatusEvent	Warning	Controller Service has stopped running.	Controller Service status has turned to stopped.
1.3.6.1.4.1.6876.100.1.0.80404	NSXTMPNodeServiceDataStoreStatusEvent	Warning	DataStore Service has stopped running.	DataStore Service status has turned to stopped.
1.3.6.1.4.1.6876.100.1.0.80405	NSXTMPNodeServiceHttpStatusEvent	Warning	HTTP Service has stopped running.	HTTP Service status has turned to stopped.
1.3.6.1.4.1.6876.100.1.0.80406	NSXTMPNodeServiceInstallUpgradeEvent	Warning	Install Upgrade Service has stopped running.	Install Upgrade Service status has turned to stopped.
1.3.6.1.4.1.6876.100.1.0.80407	NSXTMPNodeServiceLiagentStatusEvent	Warning	Liagent service has stopped running.	Liagent Service status has turned to stopped.

Table 20-12. Aria Operations for Networks Computed NSX Events (continued)

OID	Event Name	Default Severity	UI Name	Description
1.3.6.1.4.1.6876.100.1.0.80408	NSXTMPNodeServiceManagerStatusEvent	Warning	Manager Service has stopped running.	Manager Service status has turned to stopped.
1.3.6.1.4.1.6876.100.1.0.80409	NSXTMPNodeServiceMgmtPlaneBusStatusEvent	Warning	Management Plane Service has stopped running.	Management Service status has turned to stopped.
1.3.6.1.4.1.6876.100.1.0.80410	NSXTMPNodeServiceMigrationCoordinatorStatusEvent	Warning	Migration Co-ordinator Service has stopped running.	Migration Co-ordinator Service status has turned to stopped.
1.3.6.1.4.1.6876.100.1.0.80411	NSXTMPNodeServiceNodeMgmtStatusEvent	Warning	Node Management Service has stopped running.	Node Management Service status has turned to stopped.
1.3.6.1.4.1.6876.100.1.0.80412	NSXTMPNodeServiceNodeStatsStatusEvent	Warning	Node Statistics Service has stopped running.	Node Statistics Service status has turned to stopped.
1.3.6.1.4.1.6876.100.1.0.80413	NSXTMPNodeServiceNSXMessageBusStatusEvent	Warning	Message Bus Service has stopped running.	Message Bus Client Service status has turned to stopped.
1.3.6.1.4.1.6876.100.1.0.80414	NSXTMPNodeServiceNSXPlatformClientStatusEvent	Warning	Platform Client Service has stopped running.	Platform Client Service status has turned to stopped.
1.3.6.1.4.1.6876.100.1.0.80415	NSXTMPNodeServiceNSXUpgradeAgentStatusEvent	Warning	Upgrade Agent Service has stopped running.	Upgrade Service status has turned to stopped.
1.3.6.1.4.1.6876.100.1.0.80416	NSXTMPNodeServiceNTPStatusEvent	Warning	NTP Service has stopped running.	NTP Service status has turned to stopped.
1.3.6.1.4.1.6876.100.1.0.80417	NSXTMPNodeServicePolicyStatusEvent	Warning	Policy Service has stopped running.	Policy Service status has turned to stopped.
1.3.6.1.4.1.6876.100.1.0.80418	NSXTMPNodeServiceSearchStatusEvent	Warning	Search Service has stopped running.	Search Service status has turned to stopped.
1.3.6.1.4.1.6876.100.1.0.80419	NSXTMPNodeServiceSNMPStatusEvent	Warning	SNMP Service has stopped running.	SNMP Service status has turned to stopped.
1.3.6.1.4.1.6876.100.1.0.80420	NSXTMPNodeServiceSSHStatusEvent	Warning	SSH Service has stopped running.	SSH Service status has turned to stopped.
1.3.6.1.4.1.6876.100.1.0.80421	NSXTMPNodeServiceSyslogStatusEvent	Warning	Syslog Service has stopped running.	Syslog Service status has turned to stopped.

Table 20-12. Aria Operations for Networks Computed NSX Events (continued)

OID	Event Name	Default Severity	UI Name	Description
1.3.6.1.4.1.6876.100.1.0.80422	NSXTMPNodeServiceTelemetryStatusEvent	Warning	Telemetry Service has stopped running.	Telemetry Service status has turned to stopped.
1.3.6.1.4.1.6876.100.1.0.80423	NSXTMPNodeServiceUIServiceStatusEvent	Warning	UI Service has stopped running.	UI Service status has turned to stopped.
1.3.6.1.4.1.6876.100.1.0.80402	NSXTMPNodeServiceCmInventoryStatusEvent	Critical	CM Inventory Serivce has stopped	One of the Services of the NSX Management Node, namely CM Inventory Service has stopped running.
1.3.6.1.4.1.6876.100.1.0.80403	NSXTMPNodeServiceControllerStatusEvent	Critical	Controller Service has stopped	One of the Services of the NSX Management Node, namely Controller Service has stopped running.
1.3.6.1.4.1.6876.100.1.0.80404	NSXTMPNodeServiceDataStoreStatusEvent	Critical	DataStore Service has stopped	One of the Services of the NSX Management Node, namely DataStore Service has stopped running.
1.3.6.1.4.1.6876.100.1.0.80405	NSXTMPNodeServiceHttpStatusEvent	Critical	HTTP Service has stopped	One of the Services of the NSX Management Node, namely HTTP Service has stopped running.
1.3.6.1.4.1.6876.100.1.0.80406	NSXTMPNodeServiceInstallUpgradeEvent	Warning	Install Upgrade Service has stopped	One of the Services of the NSX Management Node, namely Install Upgrade Service has stopped running.
1.3.6.1.4.1.6876.100.1.0.80407	NSXTMPNodeServiceLiagentStatusEvent	Warning	Liagent service has stopped	One of the Services of the NSX Management Node, namely LI Agent Service has stopped running.
1.3.6.1.4.1.6876.100.1.0.80408	NSXTMPNodeServiceManagerStatusEvent	Critical	Manager Service has stopped	One of the Services of the NSX Management Node, namely Manager Service has stopped running.

Table 20-12. Aria Operations for Networks Computed NSX Events (continued)

OID	Event Name	Default Severity	UI Name	Description
1.3.6.1.4.1.6876.100.1.0.80409	NSXTMPNodeServiceMgmtPlaneBusStatusEvent	Warning	Management Plane Service has stopped	One of the Services of the NSX Management Node, namely Management Plane Bus Service has stopped running.
1.3.6.1.4.1.6876.100.1.0.80410	NSXTMPNodeServiceMigrationCoordinatorStatusEvent	Warning	Migration Co-ordinator Service has stopped	One of the Services of the NSX Management Node, namely Migration Co-ordinator Service has stopped running.
1.3.6.1.4.1.6876.100.1.0.80411	NSXTMPNodeServiceNodeMgmtStatusEvent	Critical	Node Management Service has stopped	One of the Services of the NSX Management Node, namely Node Management Service has stopped running.
1.3.6.1.4.1.6876.100.1.0.80412	NSXTMPNodeServiceNodeStatsStatusEvent	Critical	Node Statistics Service has stopped	One of the Services of the NSX Management Node, namely Node Statistics has stopped running.
1.3.6.1.4.1.6876.100.1.0.80413	NSXTMPNodeServiceNSXMessageBusStatusEvent	Warning	Message Bus Service has stopped	One of the Services of the NSX Management Node, namely Message Bus Service has stopped running.
1.3.6.1.4.1.6876.100.1.0.80414	NSXTMPNodeServiceNSXPlatformClientStatusEvent	Critical	Platform Client Service has stopped	One of the Services of the NSX Management Node, namely Platform Client Service has stopped running.
1.3.6.1.4.1.6876.100.1.0.80415	NSXTMPNodeServiceNSXUpgradeAgentStatusEvent	Warning	Upgrade Agent Service has stopped	One of the Services of the NSX Management Node, namely Upgrade Agent Service has stopped running.
1.3.6.1.4.1.6876.100.1.0.80416	NSXTMPNodeServiceNTPStatusEvent	Critical	NTP Service has stopped	One of the Services of the NSX Management Node, namely NTP Service has stopped running.

Table 20-12. Aria Operations for Networks Computed NSX Events (continued)

OID	Event Name	Default Severity	UI Name	Description
1.3.6.1.4.1.6876.100.1.0.80417	NSXTMPNodeServicePolicyStatusEvent	Critical	Policy Service has stopped	One of the Services of the NSX Management Node, namely Policy Service has stopped running.
1.3.6.1.4.1.6876.100.1.0.80418	NSXTMPNodeServiceSearchStatusEvent	Critical	Search Service has stopped	One of the Services of the NSX Management Node, namely Search Service has stopped running.
1.3.6.1.4.1.6876.100.1.0.80419	NSXTMPNodeServiceSNMPStatusEvent	Warning	SNMP Service has stopped	One of the Services of the NSX Management Node, namely SNMP Service has stopped running.
1.3.6.1.4.1.6876.100.1.0.80420	NSXTMPNodeServiceSSHStatusEvent	Critical	SSH Service has stopped	One of the Services of the NSX Management Node, namely SSH Service has stopped running.
1.3.6.1.4.1.6876.100.1.0.80421	NSXTMPNodeServiceSyslogStatusEvent	Critical	Syslog Service has stopped	One of the Services of the NSX Management Node, namely Syslog Service has stopped running.
1.3.6.1.4.1.6876.100.1.0.80422	NSXTMPNodeServiceTelemetryStatusEvent	Warning	Telemetry Service has stopped	One of the Services of the NSX Management Node, namely Telemetry Service has stopped running.
1.3.6.1.4.1.6876.100.1.0.80423	NSXTMPNodeServiceUIServiceStatusEvent	Critical	UI Service has stopped	One of the Services of the NSX Management Node, namely UI Service has stopped running.
1.3.6.1.4.1.6876.100.1.0.80424	NSXTMPNodeServiceClusterManagerStatusEvent	Critical	Cluster Manager Service has stopped	One of the Services of the NSX Management Node, namely Cluster Manager Service has stopped running.

Network Monitoring

21

The topics in this section show you how to configure monitoring using Internet Protocol Flow Information Export (IPFIX) profiles for the firewall and switches, as well as how to configure an IPFIX collector.

This chapter includes the following topics:

- Add an IPFIX Collector
- Add a Firewall IPFIX Profile
- Add a Switch IPFIX Profile
- IPFIX Monitoring on a vSphere Distributed Switch
- Add a Port Mirroring Session
- Port Mirroring on a vSphere Distributed Switch
- Perform a Traceflow
- Simple Network Management Protocol (SNMP)
- Network Latency Statistics
- Monitoring Tools in Manager Mode
- Checking CPU Usage and Network Latency
- Live Traffic Analysis

Add an IPFIX Collector

You can configure IPFIX collectors for firewalls and switches.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Select **Plan & Troubleshoot > IPFIX**.
3. Click the **Collectors** tab.
4. Select **Add New Collector > IPFIX Switch** or **Add New Collector > IPFIX Firewall**.
5. Enter a name.

6 Enter the IP address and port of up to four collectors. Both IPv4 and IPv6 addresses are supported.

7 Click **Save**.

Add a Firewall IPFIX Profile

You can configure IPFIX profiles for firewalls.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Plan & Troubleshoot > IPFIX**.

3 Click the **Firewall IPFIX Profiles** tab.

4 Click **Add Firewall IPFIX Profile**.

5 Complete the following details.

Setting	Description
Name and Description	Enter a name and optionally a description. **Note** If you want to create a global profile, name the profile `Global`. A global profile cannot be edited or deleted from the UI, but you can do so using NSX APIs.
Active Flow Export Timeout (Minutes)	The length of time after which a flow will time out, even if more packets associated with the flow are received. Default is 1.
Observation Domain ID	This parameter identifies which observation domain the network flows originate from. The default is 0 and indicates no specific observation domain.
Collector Configuration	Select a collector from the drop-down menu.
Applied To	Click **Set** and select a group to apply the filter to, or create a new group.
Priority	This parameter resolves conflicts when multiple profiles apply. The IPFIX exporter will use the profile with the highest priority only. A lower value means a higher priority.

6 Click **Save** and then **Yes** to continue configuring the profile.

7 Click **Save**.

Add a Switch IPFIX Profile

You can configure IPFIX profiles for switches, also known as segments.

Flow-based network monitoring enable network administrators to gain insight into traffic traversing a network.

Starting with NSX 4.0.1.1, vSphere Distributed Services Engine provides the ability to offload some of the network operations from your server CPU to a Data Processing Unit (DPU also known as SmartNIC). vSphere 8.0 supports NVIDIA BlueFiled and AMD Pensando DPU devices only.

For more information about VMware vSphere Distributed Services Engine, see *Introducing VMware vSphere® Distributed Services EngineTM and Networking Acceleration by Using DPUs* in the VMware vSphere® product documentation.

Note If you want to configure IPFIX on DPU backed VDS, you must create vmknic on 'ops' TCP/IP stack. Else, the flow information is not exported to collector.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Plan & Troubleshoot > IPFIX**.

3 Click the **Switch IPFIX Profiles** tab.

4 Click **Add Switch IPFIX Profile**.

5 Enter the following details:

Setting	Description
Name and Description	Enter a name and optionally a description. **Note** To create a global profile, name the profile `Global`. You cannot edit or delete a global profile from the UI, but you can do so using NSX APIs.
Active Timeout (seconds)	The length of time after which a flow times out, even if more packets associated with the flow are received. Default is 300.
Idle Timeout (seconds)	The length of time after which a flow times out, if no more packets associated with the flow are received. Default is 300.
Packet Sampling Probability (%)	An estimate of the percentage of packets that will be sampled. Increasing this setting can have a performance impact on the hypervisors and collectors. If all hypervisors are sending more IPFIX packets to the collector, the collector might not be able to collect all packets. Setting the probability at the default value of 0.1% decreases the performance impact.
Collector Configuration	Select a collector from the drop-down menu.
Applied To	Select a category: Segment, Segment Port, Groups, or Selected Count. The IPFIX profile applies to the selected object.
Priority	This parameter resolves conflicts when multiple profiles apply. The IPFIX exporter uses the profile with the highest priority only. A lower value means a higher priority.
Max Flows	The maximum flows cached on a bridge. Default is 16384.
Observation Domain ID	The observation domain ID identifies which observation domain the network flows originate from. Enter 0 to indicate no specific observation domain.

Setting	Description
Export Overlay Flow	This parameter defines whether to sample and export the overlay flows on uplink and tunnel ports. Both the vNIC flow and overlay flow are included in the sample. The default is **Enabled**. When disabled, only vNIC flows are sampled and exported.
Tags	Enter a tag to make searching easier.

6 Click **Save** and then **Yes** to continue configuring the profile.

7 Click **Applied To** to apply the profile to an NSGroup. You can select one or more NSGroups.

Note IPFIX Profile supports NSGroups with member types: Other NSGroups, Segment, and Segment Port. To learn more about NSGroup, see Create an NSGroup in Manager Mode.

8 Click **Save**.

IPFIX Monitoring on a vSphere Distributed Switch

Configure IPFIX monitoring for NSX Distributed Virtual port groups, and vSphere Distributed Virtual port groups that are connected to a VDS switch enabled to support NSX networking.

From vSphere, enable IPFIX for Distributed Virtual port groups (vSphere) and from NSX Manager, enable IPFIX for segments (NSX) created on a VDS switch.

Starting with NSX 4.0.1.1, vSphere Distributed Services Engine provides the ability to offload some of the network operations from your server CPU to a Data Processing Unit (DPU also known as SmartNIC). vSphere 8.0 supports NVIDIA BlueFiled and AMD Pensando DPU devices only.

For more information about VMware vSphere Distributed Services Engine, see *Introducing VMware vSphere® Distributed Services EngineTM and Networking Acceleration by Using DPUs* in the VMware vSphere® product documentation.

If you want to configure IPFIX on DPU backed VDS, you must create vmknic on 'ops' TCP/IP stack.

To enable IPFIX monitoring for Distributed Virtual port groups, see the *vSphere Networking* documentation.

To enable IPFIX monitoring for NSX port groups, see Add a Switch IPFIX Profile.

A VDS switch enabled for NSX displays the following behavior:

- Both non-uplink and uplink ports support bidirectional traffic incoming and outgoing on:
  - Ports, port groups, VMs on vSphere.
  - Segments, segment ports, and groups on NSX

- IPFIX profile samples packets on uplink ports when are coming from or going to non-uplink ports that are IPFIX enabled. For example, consider that *VM-A* and *VM-B* are connected to non-uplink ports (port-1, port-2), where port-1 connected to *VM-A* is IPFIX enabled, and port-2 connected to *VM-B* is not IPFIX enabled. When you send traffic from *VM-A* and *VM-B* traffic to port-1, only packets from *VM-A* are sampled because IPFIX is enabled only on the port that *VM-A* is connected to. IPFIX does not sample packets coming from the port-2 associated to *VM-B* because IPFIX is not enabled on that port.
- Packet count exported to the IPFIX collector is the total count based on a sampling rate, not the sampled packets. For example, IPFIX calculates the count of total packets and exports the info. For 100 incoming packets, IPFIX might sample 9–11 packets. It exports 90 or 110 packets to the IPFIX collector.

Add a Port Mirroring Session

You can use port mirroring to analyze network traffic for debugging or troubleshooting purposes. Port mirroring allows you to copy all network packets or specific packets that are seen on the segment port (or an entire segment) to another segment port.

Logical Span session type is supported only for overlay segments and not for VLAN segments.

Port mirroring is supported on ENS fast path for Remote L3 Span session type.

Note Port Mirroring is not recommended for monitoring because when used for longer durations performance is impacted.

Starting with NSX 4.0.1.1, vSphere Distributed Services Engine provides the ability to offload some of the network operations from your server CPU to a Data Processing Unit (DPU also known as SmartNIC). vSphere 8.0 supports NVIDIA BlueFiled and AMD Pensando DPU devices only.

For more information about VMware vSphere Distributed Services Engine, see *Introducing VMware vSphere® Distributed Services Engine™ and Networking Acceleration by Using DPUs* in the VMware vSphere® product documentation.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Select **Plan & Troubleshoot > Port Mirroring**.
3. Select **Add Session > Remote L3 Span** or **Add Session > Logical Span**.
4. Enter a name and optionally a description.

5 Configure the properties of the port mirroring session.

Session Type	Properties
Remote L3 Span	■ **Direction** - Select **Bidirectional**, **Ingress**, or **Egress**. ■ **TCP/IP Stack** - Select **Default** or **Mirror**. To use **Mirror**, you must bind the vmknic to the mirror stack in vSphere. **Note** If you want to configure ERSPAN on DPU backed VDS, you must create vmknic on 'mirror' TCP/IP stack. The behaviour of ERPSAN on AMD Pensando DPU and NVIDIA BlueFiled DPU is different: ■ AMD Pensando DPU device supports full offloading. This means that the mirroring is completely handled by hardware (DPU). However, AMD Pensando supports only 6 destination IPs. If it exceeds them, it will use partial offloading. ■ NVIDIA BlueFiled DPU device supports partial offloading. This means that the production traffic is handled by hardware (DPU) and mirrored packets are handled by software (ESXi on DPU). ■ **Snap Length** - Specify the number of bytes to capture from a packet. If this parameter is specified, the packet is truncated to the specified length. If not specified, the entire packet is mirrored. Supported range of values is 60–65535. ■ **Encapsulation Type** - Select **GRE**, **ERSPAN TWO**, or **ERSPAN THREE**. ■ **GRE Key** - Specify a 32-bit GRE key if encapsulation type is **GRE**. ■ **ERSPAN ID** - Specify an ERSPAN ID if encapsulation type is **ERSPAN TWO** or **ERSPAN THREE**. Supported range of values is 0–1023. The physical switch uses the ERSPAN ID to forward the mirrored traffic.
Logical Span	■ **Direction** - Select **Bidirectional**, **Ingress**, or **Egress**. ■ **Snap Length** - Specify the number of bytes to capture from a packet. If this parameter is specified, the packet is truncated to the specified length. If not specified, the entire packet is mirrored. Supported range of values is 60–65535.

6 Click **Set** in the **Source** column to set a source.

For a Logical Span session, the available sources are:

- Segment port
- Group of virtual machines
- Group of virtual network interfaces

For a Remote L3 Span session, the available sources are:

- Segment
- Segment port
- Group of virtual machines
- Group of virtual network interfaces

The following restrictions apply when you select a group of VMs or a group of virtual network interfaces:

- The group can have a maximum of six VMs that are statically added.

- The group can have a maximum of six virtual network interfaces that are statically added.

7 Click **Set** in the **Destination** column to set a destination.

For a Logical Span session, the available destinations are:

- Group of virtual machines
- Group of virtual network interfaces

The following restrictions apply when you select a group of VMs or a group of virtual network interfaces:

- The group can have a maximum of three VMs that are statically added.
- The group can have a maximum of three virtual network interfaces that are statically added.

For a Remote L3 Span session, the available destination is an IP Addresses Only group. The group can have a maximum of three IPs.

8 (Optional) Instead of mirroring all the network packets from the source, you can filter the packets that are captured for port mirroring.

a Expand the **Advanced Mirroring Filters** section.

b Select an **Action**.

Action	Description
Include	Packets that match the filter are mirrored.
Exclude	Packets that do not match the filter are mirrored.

c Next to **Filters**, click **Set**, and then click **Add Filter**.

d Specify the filter properties.

Only one filter is supported.

Property	Description
Protocol	The transport protocol that is used to filter the packets. Available options are TCP, UDP.
Source IPs	The source IP address, IP range, or IP prefix that is used to filter the packets.
Source Port	The source port or port range that is used to filter the packets.
Destination IPs	The destination IP address, IP range, or IP prefix that is used to filter the packets.
Destination Port	The destination port or port range that is used to filter the packets.

9 Click **Save**.

Port Mirroring on a vSphere Distributed Switch

You can configure port mirroring for port groups, virtual NICs of VMs, and VMs created in NSX and vSphere Distributed Virtual port groups created in vSphere that are connected to a vSphere Distributed Switch (VDS) switch.

In vCenter Server, configure port mirroring for vSphere Distributed Virtual port groups on a VDS switch.

In NSX Manager, configure port mirroring for segments (in NSX) on a VDS switch.

Note You cannot edit configuration for a segment created in NSX in vCenter Server. As an admin, you can view the properties of a port mirroring session to know on which switch it is created.

Starting with NSX 4.0.1.1, vSphere Distributed Services Engine provides the ability to offload some of the network operations from your server CPU to a Data Processing Unit (DPU also known as SmartNIC). vSphere 8.0 supports NVIDIA BlueFiled and AMD Pensando DPU devices only.

For more information about VMware vSphere Distributed Services Engine, see *Introducing VMware vSphere® Distributed Services EngineTM and Networking Acceleration by Using DPUs* in the VMware vSphere® product documentation.

If you want to configure ERSPAN on DPU backed VDS, you must create vmknic on 'mirror' TCP/IP stack.

To enable port mirroring on vSphere Distributed Virtual port groups, see the *vSphere Networking* documentation.

To enable port mirroring on segments, ports, groups in NSX from both Policy and Manager modes in NSX Manager, see:

- Add a Port Mirroring Session
- Monitor Port Mirroring Sessions in Manager Mode

Uplink Conflict Between Teaming and Remote SPAN

In vSphere, by default, the **Remote SPAN** in a teaming policy is set to `Disallowed`. If you use all the available physical NICs to configure remote SPAN, there are no free uplinks available for the teaming policy to consume. The unavailability of any free uplink means that uplink traffic is not allowed on destination ports, resulting in configuration errors.

However, in NSX, by default, **Normal I/O on Destination Ports** is set to `Allowed`. In NSX, port mirroring configured for NSX port groups on VDS switch allows teaming and port mirroring on destination ports. So, uplink configuration errors do not occur in NSX.

To resolve a uplink conflict when configuring teaming and remote SPAN:

- Ensure that a free uplink is available. For example, on an ESXi host with 2 physical NICs, do not assign both these uplinks as destination IP addresses in the remote span port mirroring profile to avoid uplink conflicts in configuration. There must be at least one available uplink that can be configured in teaming profile.
- In vCenter Server, edit the port mirror configuration profile and set the **Normal I/O on Destination Ports** to **`Allowed`**.

Perform a Traceflow

Use Traceflow to inspect the path of a packet. Traceflow traces the transport node-level path of a packet. The trace packet traverses the segment overlay, but is not visible to interfaces attached to the segment. In other words, no packet is actually delivered to the test packet's intended recipients.

For a VLAN-backed segment, enable In-band Network Telemetry (INT) by calling the `PUT /api/v1/infra/ops-global-config` API.

Sample request:

URL:

```
PUT https:/{{nsx-manager-ip}}/policy/api/v1/infra/ops-global-config
```

Body

```
{
 "display_name": "ops-global-config",
 "in_band_network_telementry": {
 "dscp_value": 2,
 "indicator_type": "DSCP_VALUE"
 },
"path": "/infra/ops-global-config",
"relative_path": "ops-global-config",
 "_revision": 0
}
```

INT cannot be configured if a traceflow request is already in progress. For more information about APIs for In-band Network Telemetry, see *NSX API Guide*. VLAN traceflow supports only TCP/UDP and ICMP packets.

VLAN tracing is not supported on Edge nodes. Attempting to inject a trace packet on an Edge node results in an API validation error. For a VLAN traceflow, if the injected trace packet traverses an Edge node, the resultant trace is incomplete and only observations on ESX nodes are displayed.

From 4.0 or later, you can view the IPSec VPN specific observations when the packet is processed.

Note NSX Traceflow does not work with HCX extended networks.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Plan & Troubleshoot > Traffic Analysis > Traceflow > Get Started**.

3 Select an IPv4 or IPv6 address type.

4 Select a traffic type.

For IPv4 addresses the traffic type choices are Unicast, Multicast, and Broadcast. For IPv6 address the traffic type choices are Unicast or Multicast.

Note: Multicast and broadcast are not supported in a VMware Cloud (VMC) environment.

5 (Optional) Select a protocol and provide related information.

Protocol	Parameters
DHCP	Select a DHCP OP code: **Boot Request** or **Boot Reply**.
DHCPv6	Select a DHCP message type: **Solicit**, **Advertise**, **Request**, or **Reply**. **Note** This option is available only when IPv6 is selected for IP Address.
DNS	Specify an address and select a message type: **Query** or **Response**.
ICMP	Specify an ICMP ID and a sequence.
ICMPv6	Specify an ICMP ID and a sequence. **Note** This option is available only when IPv6 is selected for IP Address.
TCP	Specify a source port, a destination port, and TCP flags.
UDP	Specify a source port and a destination port.
ARP	Select an ARP OP Code: **ARP Request** or **ARP Reply** .
NDP	For Unicast traffic, specify the **Destination IP** address. For Multicast traffic, specify the **Destination IP** and **Destination MAC** address. **Note** This option is available only when IPv6 is selected for IP Address.

For the TCP protocol, note the following:

- The default flag is SYN.
- SYN cannot be combined with RST or FIN.
- If SYN is not selected, you must select ACK or RST.
- ACK cannot be combined with FIN, PSH, or URG.

6 Specify the source and destination information according to the traffic type.

Note For VLAN tracing across different VLAN segments, ensure that you set the appropriate MAC address using one of the following options:

- In the NSX Manager user interface, select **IP – Mac** as destination type and select **Layer 2**. Type the destination IP address and the MAC address of the physical gateway that can forward the packet to the destination VLAN network..
- When using APIs to perform a traceflow, set `routed` as `false` and use the MAC address of the physical gateway that can forward the packet to the destination VLAN network. For more information about traceflow APIs, see *NSX API Guide*.

Traffic Type	Source	Destination
Unicast	Select a VM or a logical port. For a VM: ■ Select a VM from the drop-down list. ■ Select a virtual interface. ■ The IP address and MAC address are displayed if VMtools is installed in the VM, or if the VM is deployed using OpenStack plug-in (address bindings will be used in this case). If the VM has more than one IP address, select one from the drop-down list. ■ If the IP address and MAC address are not displayed, enter the IP address and MAC address in the text boxes. For a logical port: ■ Select an attachment type: **VIF**, **DHCP**, **Edge Uplink**, or **Edge Centralized Service**. ■ Select a port.	Select a VM, a logical port, or IP-MAC. For a VM: ■ Select a VM from the drop-down list. ■ Select a virtual interface. ■ The IP address and MAC address are displayed if VMtools is installed in the VM or if the VM is deployed using OpenStack plug-in (address bindings will be used in this case). If the VM has more than one IP address, select one from the drop-down list. ■ If the IP address and MAC address are not displayed, enter the IP address and MAC address in the text boxes. For a logical port: ■ Select an attachment type: **VIF**, **DHCP**, **Edge Uplink**, or **Edge Centralized Service**. ■ Select a port. For IP-MAC: ■ Select the trace type (layer 2 or layer 3). For layer 2, enter an IP address and a MAC address. For layer 3, enter an IP address.
Multicast	Same as above.	Enter an IP Address. It must be a multicast address from 224.0.0.0 - 239.255.255.255.
Broadcast	Same as above.	Enter a subnet prefix length.

7 (Optional) Click **Advanced Settings** to see the advanced options.

In the left column, enter the desired values or input for the following fields:

Option	Description
Frame Size	The default is 128.
TTL	The default is 64.

Option	Description
Timeout (ms)	The default is 10000.
Ethertype	The default is 2048.
Payload Type	Select **Base64**, **Hex**, **Plaintext**, **Binary**, or **Decimal**.
Payload Data	Payload formatted based on selected type.

8 Click **Trace**.

The output includes a graphical map of the topology and a table listing the observed packets. The first packet listed has the observation type `Injected` and shows the packet that is injected at the injection point.

You can apply a filter (**All**, **Delivered**, **Dropped**) on the observations that are displayed. If there are dropped observations, the **Dropped** filter is applied by default. Otherwise, the **All** filter is applied.

The graphical map shows the backplane and router links. Note that bridging information is not displayed.

Simple Network Management Protocol (SNMP)

You can use Simple Network Management Protocol (SNMP) to monitor your NSX components. The SNMP service is not started by default after installation.

The SNMP Framework in NSX enables you to monitor various system entities (such as disk on NSX Edge) and logical entities (such as NSX Edge VPN tunnel) using their SNMP managers. This framework enables NSX verticals and platform to define SNMP MIB objects to be monitored and which can be used to enable their SMNP managers to interact with NSX.

To download the SNMP MIB files, see Knowledge Base article 1013445: SNMP MIB module file download. For NSX, download the folder and use the extracted file **VMWARE-NSX-MIB.mib**.

For SNMP configuration, see *Configure SNMP for ESXi* in the VMware vSphere product documentation.

Procedure

1 Log in to the NSX Manager CLI or the NSX Edge CLI.

2 Run the following commands

- For SNMPv1/SNMPv2:

```
set snmp community <community-string>
start service snmp
```

The maximum character limit for **`community-string`** is 64.

- For SNMPv3

```
set snmp v3-users <user_name> auth-password <auth_password> priv-password
<priv_password>

start service snmp
```

The maximum character limit for **user_name** is 32. Ensure that your passwords meet PAM constraints. If you want to change the default engine id, use the following command:

```
set snmp v3-engine-id <v3-engine-id>

start service snmp
```

v3-engine-id is an even-length hexadecimal string that is 10 to 64 characters long and cannot be all 0s or Fs.

NSX supports SHA1 and AES128 as the authentication and privacy protocols. You can also use API calls to set up SNMPv3. For more information, see the *NSX API Guide*.

Network Latency Statistics

In a network, latency can accumulate at multiple endpoints in the data path. As a network administrator, you need the ability to monitor the latency of a network to diagnose and troubleshoot performance bottlenecks in the network.

The following network latency statistics can be measured on host transport nodes:

- pNIC to vNIC
- vNIC to pNIC
- vNIC to vNIC
- VTEP to VTEP

In NSX, the following limitations apply to measuring latency statistics:

- Only ESXi host transport nodes are supported for measuring network latency in the data plane.
- Edge transport nodes are not supported.
- On VLAN segments, network latency is measured only when the two vNICs belong to VMs on the same ESXi host.
- When the VMs are attached to separate segments, network latency is measured only when the data traffic is routed through the distributed router (DR) instance on the ESXi host transport nodes. If the data traffic is routed through the DR instance on the edge transport nodes, network latency is not measured.
- Enhanced networking stack (ENS) does not support vNIC-to-pNIC, pNIC-to-vNIC, or vNIC-to-vNIC latency.

- Latency measurement is not supported when an east-west network traffic protection is configured using partner service VMs. Latency monitoring is disabled on the ports of service virtual machine (SVMs) and guest VMs.
- Latency measurement is not supported on the Data Processing Unit (DPU).

You can export the latency data to external network performance monitoring tools and run analytics on the data. The external monitoring tools are also called collectors. By using a collector, you can achieve greater network visibility, optimize network performance, and identify the endpoints in the data path that cause a significant latency in the network.

After the hosts are configured to measure network latency statistics, the network operations agent (netopa) on the hosts periodically polls the data plane. When latency data is available, the agent exports the data at preconfigured intervals to the external collectors.

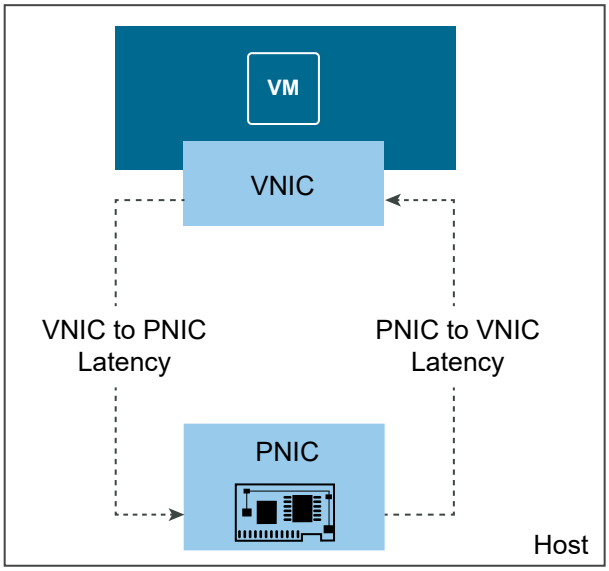
Note

- The netopa agent can export the network latency statistics only to Aria Operations for Networks. Other collector tools are not supported currently.
- For VTEP statistics, Aria Operations for Networks 5.0 or later is required. For the pNIC and vNIC statistics, Aria Operations for Networks 5.3 or later is required.
- You can configure ESXi hosts to measure network latency statistics only by using the NSX REST APIs.

You can measure network latency statistics for both standalone ESXi hosts and ESXi hosts that are a part of the VMware vCenter cluster. However, Aria Operations for Networks supports collecting latency statistics only from vCenter-managed ESXi hosts and not from standalone ESXi hosts.

pNIC-to-vNIC and vNIC-to-pNIC Latency

When pNIC latency measurement is enabled on a host transport node, vNIC-to-pNIC latency and pNIC-to-vNIC latency are computed for each vNIC on the host transport node.



pNIC-to-vNIC and vNIC-to-pNIC latency statistics are exported to the external collector in the following format:

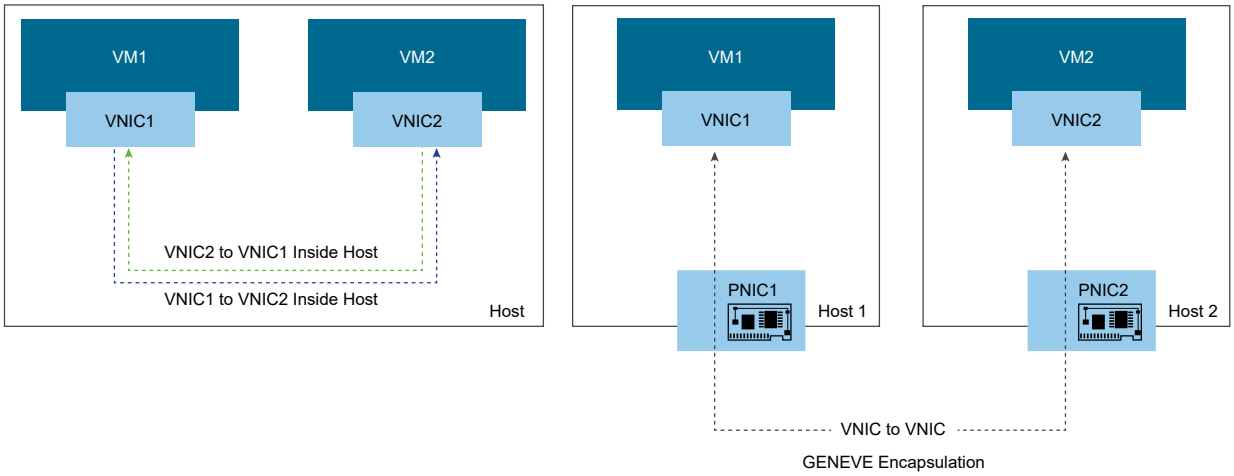
Endpoint1, Endpoint2, Max, Min, Avg

Where:

- *Endpoint1, Endpoint2* can either be the virtual interface ID (VIF ID) or the name of the physical adapter on an ESXi host (vmnic).
- *Max*, *Min*, and *Avg* represent the maximum, minimum, and average latency values between the two endpoints in microseconds.

vNIC to vNIC Latency

This latency represents the time taken by the data packet to travel from the source vNIC to the destination vNIC either on the same ESXi host or different ESXi hosts. If the vNICs are on different ESXi hosts, only GENEVE encapsulation protocol is supported in the overlay tunnel between the hosts.



vNIC to vNIC network latency is computed as follows:

- When the source VNIC1 on VM1 and the destination VNIC2 on VM2 are on the same host, a single-trip latency is calculated for each trip and exported to the collector. In other words, latency for each trip from VNIC1 to VNIC2 and from VNIC2 to VNIC1 is computed separately.
- When the source VNIC1 on VM1 and the destination VNIC2 on VM2 are on different hosts, total round-trip latency is calculated, and only a single latency value is exported to the collector. If there is no return traffic from VNIC2 to VNIC1, no network latency is exported to the collector.

Note NSX calculates the vNIC-to-vNIC latency between hosts directly by using the timestamps in the GENEVE encapsulated packets. You do not have to enable pNIC latency measurement on the host and the VTEP-to-VTEP latency. The pNIC-to-vNIC, vNIC-to-pNIC, and VTEP-to-VTEP statistics are independent of the vNIC-to-vNIC statistic.

vNIC-to-vNIC latency statistics are exported to the external collector in the following format:

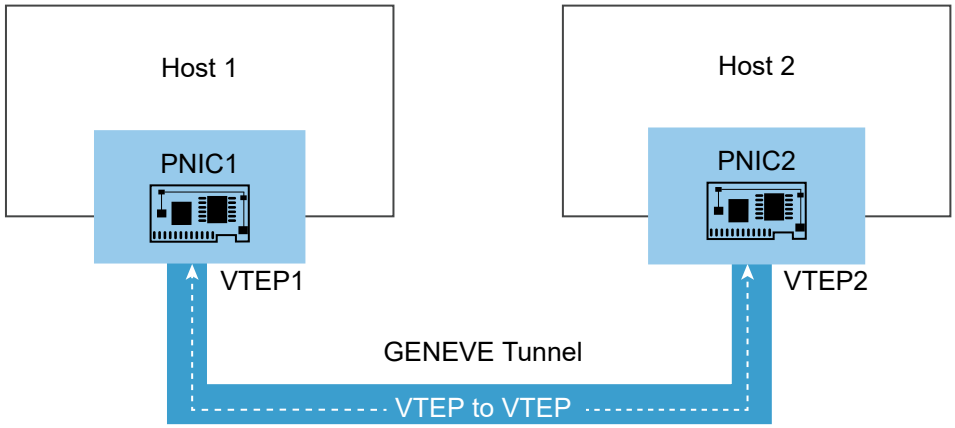
VIF1, VIF2, Max, Min, Avg

Where:

- *VIF1, VIF2* represent the virtual interfaces or the vNICs.
- *Max*, *Min*, and *Avg* represent the maximum, minimum, and average latency values between the two vNICs in microseconds.

VTEP-to-VTEP Latency

This latency represents the total round-trip time taken by the data packet to travel from the source VTEP to the destination VTEP. To measure VTEP-to-VTEP latency, you must enable latency in the transport zone profile.



To calculate the VTEP-to-VTEP latency between ESXi hosts, Bidirectional Flow Detection (BFD) protocol is used. NSX extends the BFD protocol with timestamps in the payload to support latency computation between the VTEPs. BFD packets are transmitted at regular intervals in each tunnel between the hosts to compute the VTEP-to-VTEP latency.

- Measure Network Latency Statistics

  You can configure ESXi hosts in your network to measure: pNIC to vNIC, vNIC to pNIC, vNIC to vNIC, and VTEP to VTEP network latency statistics.

- Export Network Latency Statistics

  You can export network latency statistics to external collectors and run analytics on the data. The netopa agent that is running in the ESXi hosts can export the network latency statistics only to Aria Operations for Networks. Other collector tools are not supported currently.

Measure Network Latency Statistics

You can configure ESXi hosts in your network to measure: pNIC to vNIC, vNIC to pNIC, vNIC to vNIC, and VTEP to VTEP network latency statistics.

Configuration is supported only using the NSX REST APIs. The steps in the following procedure list the NSX Policy APIs that you must run to configure the calculation of various network latency statistics. For a detailed information about the API schema, example request, example response, and error messages of all the APIs, you must read the *NSX API Guide*.

Prerequisites

Both vCenter-managed hosts and standalone ESXi hosts that you want to configure for measuring network latency statistics must be prepared for NSX. That is, NSX components must be installed on all the ESXi hosts in your network.

Procedure

1 To compute vNIC to vNIC, pNIC to vNIC, and vNIC to pNIC network latency statistics, do these steps:

a Create a group that contains host transport nodes as static members by using the following PUT API:

```
PUT https://<nsx-mgr>/policy/api/v1/infra/domains/<domain-id>/groups/<group-id>
```

In the request payload of this PUT API, specify the host transport node IDs in the `expression` parameter, as shown in the following example:

Example PUT Request

```
PUT https://<nsx-mgr>/policy/api/v1/infra/domains/default/groups/TNGroup

{
    "expression": [
        {
            "paths": [
                "/infra/sites/default/enforcement-points/default/host-transport-nodes/
4efdb573-fcce-43ff-8b35-dac583a86239"
            ],
            "resource_type": "PathExpression"
        }
    ],
    "extended_expression": [],
    "reference": false,
    "group_type": [],
    "resource_type": "Group",
    "id": "TNGroup",
    "display_name": "TNGroup",
    "path": "/infra/domains/default/groups/TNGroup",
    "relative_path": "TNGroup",
    "parent_path": "/infra/domains/default"
}
```

Observe that in this example request, the `expression` parameter contains a single host transport node ID.

b Create a latency profile with the following PUT API:

```
PUT https://<nsx-mgr>/policy/api/v1/infra/latency-profiles/<profile-id>
```

By default, vNIC to vNIC latency is measured for all the vNICs on the host transport node.

In the request body of this API, configure the following information:

- Activate or deactivate pNIC latency on the host. When it is activated, pNIC to vNIC and vNIC to pNIC latency are calculated for each vNIC on the host transport node.
- Specify either the sampling rate or the sampling interval.

If you configure both, sampling interval takes a precedence.

- Specify the path to the group that you created in the earlier step.

Example PUT Request

```
PUT https://<nsx-mgr>/policy/api/v1/infra/latency-profiles/profile1
    {
      "sampling_rate": 100,
      "sampling_interval": 10,
      "pnic_latency_enabled": false,
      "applied_to_group_path": "/infra/domains/default/groups/TNGroup"
    }
```

2 To measure VTEP to VTEP latency statistics, enable latency in the BFD health monitoring profile, which is a resource type in the transport zone profile. Run the following PUT API:

```
PUT https://<nsx-mgr>/policy/api/v1/infra/transport-zone-profiles/<tz-profile-id>
```

What to do next

Export the statistics to an external collector for a deeper network insight and troubleshooting network-specific latency problems.

Export Network Latency Statistics

You can export network latency statistics to external collectors and run analytics on the data. The netopa agent that is running in the ESXi hosts can export the network latency statistics only to Aria Operations for Networks. Other collector tools are not supported currently.

You can collect network latency statistics from only vCenter-managed ESXi hosts. Collecting latency statistics from standalone ESXi hosts is not supported.

You can export network latency statistics by using any one of the following methods:

- Method 1: Use the management plane APIs in NSX.
- Method 2: Enable an optional setting in the Aria Operations for Networks UI to collect latency statistics.

Prerequisites

- In the Aria Operations for Networks UI, complete the following tasks in the given order:
  a Add VMware vCenter as the data source. If you have multiple vCenter Servers added as Compute Managers in your NSX environment, you can add all vCenter Servers as the data source.
  b Add NSX Manager as the data source.

  For detailed information about adding data sources in Aria Operations for Networks, see the Aria Operations for Networks documentation.

- Ensure that port 1991 is open on the collector to receive network latency data from the ESXi hosts.

Procedure

1 Method 1: Use the NSX REST APIs.

   a Ensure that you have configured the ESXi hosts to measure network latency statistics.

   For detailed steps, see Measure Network Latency Statistics.

   b Export the network latency statistics to the collector with the following PUT API:

   ```
   PUT https://<manager-ip>/api/v1/global-configs/OperationCollectorGlobalConfig
   -d '<content>'
   ```

   In the request body of this API, configure the following information:

   - Details of external collectors, such as collector IP address, collector port.
   - Report interval that controls the frequency at which the netopa agent sends statistics to the collector.

2 Method 2: Enable an optional setting in the Aria Operations for Networks UI to collect latency statistics.

   When you add NSX Manager as the data source in Aria Operations for Networks, select the **Enable latency metric collection** check box. This option enables Aria Operations for Networks to collect latency statistics from the ESXi hosts.

   For detailed information about adding NSX Manager as the data source in Aria Operations for Networks, see the Aria Operations for Networks documentation.

Results

vNIC-to-vNIC latency statistics are exported to the external collector in the following format:

VIF1, VIF2, Max, Min, Avg

Where:

- *VIF1, VIF2* represent the virtual interfaces or the vNICs.
- *Max*, *Min*, and *Avg* represent the maximum, minimum, and average time between the two vNICs in microseconds.

pNIC-to-vNIC and vNIC-to-pNIC latency statistics are exported to the external collector in the following format:

Endpoint1, Endpoint2, Max, Min, Avg

Where:

- *Endpoint1, Endpoint2* can either be the virtual interface ID (VIF ID) or the name of the physical adapter on an ESXi host (vmnic).

- *Max*, *Min*, and *Avg* represent the maximum, minimum, and average time between the two endpoints in microseconds.

Monitoring Tools in Manager Mode

NSX support monitoring methods in **Manager** mode, including viewing port connections, traceflow, port mirroring, and activity monitoring.

View Port Connection Information in Manager Mode

You can use the port connection tool to quickly visualize and troubleshoot the connection between two VMs.

Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Select **Plan and Troubleshoot > Port Connection** from the navigation panel.
3. Select a VM from the **Source Virtual Machine** drop-down menu.
4. Select a VM from the **Destination Virtual Machine** drop-down menu.
5. Click **Go**.

   A visual map of the port connection topology is displayed. You can click on any of the components in the visual output to reveal more information about that component.

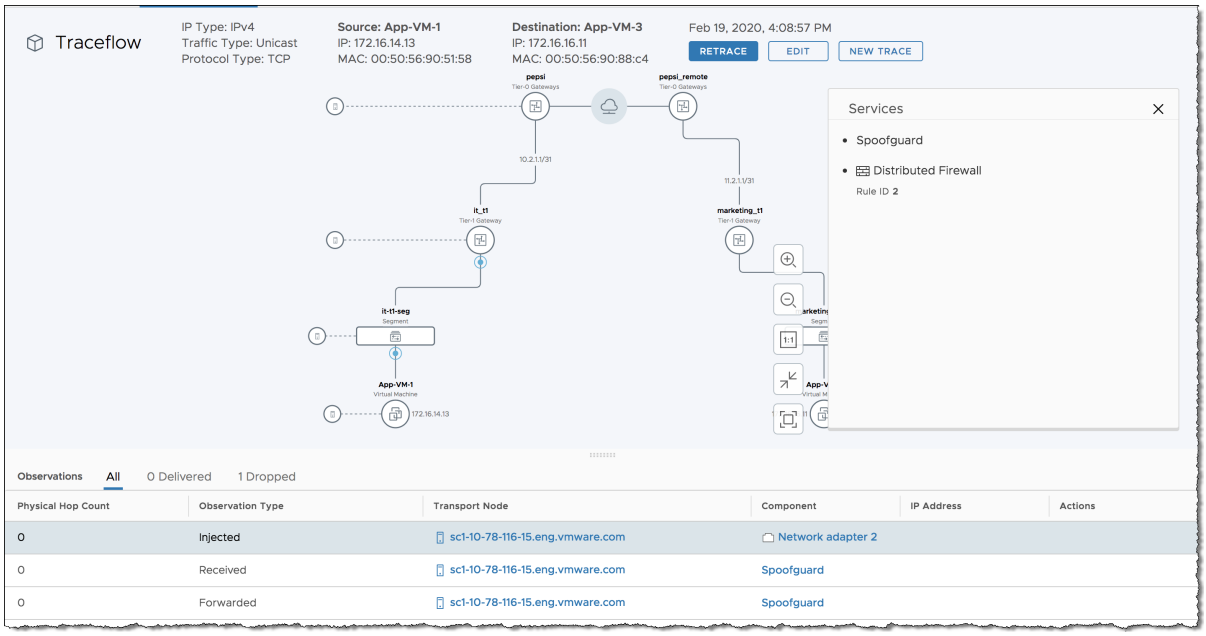
Traceflow

Traceflow allows you to inject a packet into the network and monitor its flow across the network. This flow allows you to monitor your network and identify issues such as bottlenecks or disruptions.

Traceflow allows you to identify the path (or paths) a packet takes to reach its destination or, conversely, where a packet is dropped along the way. Each entity reports the packet handling on input and output, so you can determine whether issues occur when receiving a packet or when forwarding the packet.

The NSX Manager interface graphically displays the trace route based on the parameters you set (IP address type, traffic type, source, and destination). This display page also enables you to edit the parameters, retrace the traceflow, or create a new one.

Figure 21-1. Sample traceflow diagram



What is Traceflow?

Traceflow is not the same as a ping request/response that goes from guest-VM stack to guest-VM stack. Traceflow observes a marked packet as it traverses the overlay network, and each packet is monitored as it crosses the overlay network until it reaches a destination guest VM or an Edge uplink. Note that the injected marked packet is never actually delivered to the destination guest VM.

Traceflow can be used on transport nodes and supports both IPV4 and IPv6 protocols including: ICMP, TCP, UDP, DHCP, DNS and ARP/NDP.

Traceflow Parameters

You can construct packets with custom header fields and packet sizes. The source or destination for the traceflow can be a logical switch port, logical router uplink port, CSP or DHCP port. The destination endpoint can be any device in the NSX overlay or in the underlay. However, you cannot select a destination that is north of an NSX Edge node. The destination must be on the same subnet, or must be reachable through NSX distributed logical routers.

If NSX bridging is configured, packets with unknown destination MAC addresses are always sent to the bridge. Typically, the bridge forwards these packets to a VLAN and reports the traceflow packet as delivered. A packet reported as delivered does not necessarily mean that the trace packet was delivered to the specified destination.

For a unicast traceflow packet, you can observe packet replication and/or flooding in traceflow observations.

- A traceflow packet is replicated if the logical switch does not know the TEP(s) to which the packet is destined.

- A traceflow packet is flooded if VDS does not know the virtual switch port(s) to which the packet is destined.

You can specify multicast and broadcast packets as traceflow packets.

- For multicast traffic, the source is a VM vNIC or a logical port, and the destination is a multicast IP address.
- For broadcast traffic, the source is a VM vNIC or a logical port, and the Layer 2 destination MAC address is FF:FF:FF:FF:FF:FF.

To create a valid packet for firewall inspection, the broadcast traceflow operation requires a subnet prefix length. The subnet mask enables NSX to calculate an IP network address for the packet. A multicast or broadcast traceflow packet can be delivered to multiple VM vNICs or Edge uplinks, resulting in the generation of multiple delivered observations.

Trace the Path of a Packet with Traceflow in Manager Mode

Use Traceflow to inspect the path of a packet. Traceflow traces the transport node-level path of a packet. The trace packet traverses the logical switch overlay, but is not visible to interfaces attached to the logical switch. In other words, no packet is actually delivered to the test packet's intended recipients.

Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Select **Plan & Troubleshoot > Traceflow**.
3. Select an IPv4 or IPv6 address type.
4. Select a traffic type.

   For IPv4 addresses the traffic type choices are Unicast, Multicast, and Broadcast. For IPv6 address the traffic type choices are Unicast or Multicast.

   Note: Multicast and broadcast are not supported in a VMware Cloud (VMC) environment.

5 Specify the source and destination information according to the traffic type.

Traffic Type	Source	Destination
Unicast	Select a VM or a logical port. For a VM: ■ Select a VM from the drop-down list. ■ Select a virtual interface. ■ The IP address and MAC address are displayed if VMtools is installed in the VM, or if the VM is deployed using OpenStack plug-in (address bindings will be used in this case). If the VM has more than one IP address, select one from the drop-down list. ■ If the IP address and MAC address are not displayed, enter the IP address and MAC address in the text boxes. For a logical port: ■ Select an attachment type: **VIF**, **DHCP**, **Edge Uplink**, or **Edge Centralized Service**. ■ Select a port.	Select a VM, a logical port, or IP-MAC. For a VM: ■ Select a VM from the drop-down list. ■ Select a virtual interface. ■ The IP address and MAC address are displayed if VMtools is installed in the VM or if the VM is deployed using OpenStack plug-in (address bindings will be used in this case). If the VM has more than one IP address, select one from the drop-down list. ■ If the IP address and MAC address are not displayed, enter the IP address and MAC address in the text boxes. For a logical port: ■ Select an attachment type: **VIF**, **DHCP**, **Edge Uplink**, or **Edge Centralized Service**. ■ Select a port. For IP-MAC: ■ Select the trace type (layer 2 or layer 3). For layer 2, enter an IP address and a MAC address. For layer 3, enter an IP address.
Multicast	Same as above.	Enter an IP Address. It must be a multicast address from 224.0.0.0 - 239.255.255.255.
Broadcast	Same as above.	Enter a subnet prefix length.

6 (Optional) Click **Advanced** to see the advanced options.

7 (Optional) In the left column, enter the desired values or input for the following fields:

Option	Description
Frame Size	The default is 128.
TTL	The default is 64.
Timeout (ms)	The default is 10000.
Ethertype	The default is 2048.
Payload Type	Select **Base64**, **Hex**, **Plaintext**, **Binary**, or **Decimal**.
Payload Data	Payload formatted based on selected type.

8 (Optional) Select a protocol and provide related information.

Protocol	Parameters
TCP	Specify a source port, a destination port, and TCP flags.
UDP	Specify a source port and a destination port.
ICMPv6	Specify an ICMP ID and a sequence.
ICMP	Specify an ICMP ID and a sequence.
DHCPv6	Select a DHCP message type: **Solicit**, **Advertise**, **Request**, or **Reply**.
DHCP	Select a DHCP OP code: **Boot Request** or **Boot Reply**.
DNS	Specify an address and select a message type: **Query** or **Response**.

9 Click **Trace**.

Information about the connections, components, and layers is displayed. The output includes a table listing Observation Type (Delivered, Dropped, Received, Forwarded), Transport Node, and Component, and a graphical map of the topology if unicast and logical switch as a destination are selected. You can apply a filter (**All**, **Delivered**, **Dropped**) on the observations that are displayed. If there are dropped observations, the **Dropped** filter is applied by default. Otherwise, the **All** filter is applied. The graphical map shows the backplane and router links. Note that bridging information is not displayed.

Monitor Port Mirroring Sessions in Manager Mode

You can monitor port mirroring sessions for troubleshooting and other purposes.

Note that logical SPAN is supported for overlay logical switches only and not VLAN logical switches.

NSX Cloud Note If using NSX Cloud, see NSX Features Supported with NSX Cloud for a list of auto-generated logical entities, supported features, and configurations required for NSX Cloud.

This feature has the following restrictions:

- A source mirror port cannot be in more than one mirror session.
- For a local SPAN session, the mirror session source and destination ports must be on the same host vSwitch. Therefore, if you vMotion the VM that has the source or destination port to another host, traffic on that port can no longer be mirrored.
- On ESXi, when mirroring is enabled on the uplink, raw production TCP packets are encapsulated using the Geneve protocol by VDL2 into UDP packets. A physical NIC that supports TSO (TCP segmentation offload) can change the packets and mark the packets with the MUST_TSO flag. On a monitor VM with VMXNET3 or E1000 vNICs, the driver treats the packets as regular UDP packets and cannot handle the MUST_TSO flag, and will drop the packets.

If a lot of traffic is mirrored to a monitor VM, there is a potential for the driver's buffer ring to become full and packets to be dropped. To alleviate the problem, you can take one or more of the following actions:

- Increase the rx buffer ring size.
- Assign more CPU resources to the VM.
- Use the Data Plane Development Kit (DPDK) to improve packet processing performance.

Note Make sure that the monitor VM's MTU setting is large enough to handle the packets. This is especially important for encapsulated packets because encapsulation increases the size of packets. Otherwise, packets might be dropped. This is not an issue with ESXi VMs with VMXNET3 NICs, but is a potential issue with other types of NICs.

Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Select **Plan & Troubleshoot > Port Mirroring > Port Mirroring Session**.
3. Click **Add** and select a session type.

   The available types are **Local SPAN**, **Remote SPAN**, **Remote L3 SPAN**, and **Logical SPAN**.
4. Enter a session name and optionally a description.
5. Provide additional parameters.

Session Type	Parameters
Local SPAN	■ **Transport Node** - Select a transport node. ■ **Direction** - Select **Bidirectional**, **Ingress**, or **Egress**. ■ **Packet Truncation** - Select a packet truncation value.
Remote SPAN	■ **Session Type** - Select **RSPAN Source session** or **RSPAN Destination session**. ■ **Transport Node** - Select a transport node. ■ **Direction** - Select **Bidirectional**, **Ingress**, or **Egress**. ■ **Packet Truncation** - Select a packet truncation value. ■ **Encap. VLAN ID** - Specify an encapsulation VLAN ID. ■ **Preserve Orig. VLAN** - Select whether to preserve the original VLAN ID.

Session Type	Parameters
Remote L3 SPAN	■ **Encapsulation** - Select **GRE**, **ERSPAN TWO**, or **ERSPAN THREE**. ■ **GRE Key** - Specify a GRE key if encapsulation is **GRE**. **ERSPAN ID** - Specify an ERSPAN ID if encapsulation is **ERSPAN TWO** or **ERSPAN THREE**. ■ **Direction** - Select **Bidirectional**, **Ingress**, or **Egress**. ■ **Packet Truncation** - Select a packet truncation value.
Logical SPAN	■ **Logical Switch** - Select a logical switch. ■ **Direction** - Select **Bidirectional**, **Ingress**, or **Egress**. ■ **Packet Truncation** - Select a packet truncation value.

6 Click **Next**.

7 Provide source information.

Session Type	Parameters
Local SPAN	■ Select a VDS. ■ Select physical interfaces. ■ Enable or disable encapsulated packet. ■ Select virtual machines. ■ Select virtual interfaces.
Remote SPAN	■ Select virtual machines. ■ Select virtual interfaces.
Remote L3 SPAN	■ Select virtual machines. ■ Select virtual interfaces. ■ Select a logical switch.
Logical SPAN	■ Select logical ports.

8 Click **Next**.

9 Provide destination information.

Session Type	Parameters
Local SPAN	■ Select virtual machines. ■ Select virtual interfaces.
Remote SPAN	■ Select a VDS. ■ Select physical interfaces.
Remote L3 SPAN	■ Specify an IPv4 address.
Logical SPAN	■ Select logical ports.

10 Click **Save**.

You cannot change the source or destination after saving the port mirroring session.

Configure Filters for a Port Mirroring Session

You can configure filters for port mirroring sessions to limit the amount of data that is mirrored.

This feature has the following capabilities and restrictions:

- Only ESXi host transport nodes are supported.
- IP address, IP prefix, and IP ranges are supported for source and destination.
- IPSet for source or destination is not supported.
- Mirror statistics are not supported.

You must configure filters using the API. Using the NSX Manager UI is not supported. For more information about the port mirroring API and the `PortMirroringFilter` schema, see the *NSX API Guide*.

Procedure

1. Configure a port mirroring session using the NSX Manager UI or API.
2. Call the `GET /api/v1/mirror-sessions` API to get information about the port mirroring session.
3. Call the `PUT /api/v1/mirror-sessions/<mirror-session-id>` API to add one or more filters. For example,

```
PUT https://<nsx-mgr>/api/v1/mirror-sessions/e57e8b2d-3047-4550-b230-dd1ee0e10b49
{
  "resource_type": "PortMirroringSession",
  "id": "e57e8b2d-3047-4550-b230-dd1ee0e10b49",
  "display_name": "port-mirror-session-1",
  "description": "Pnic port mirror session 1",
  "mirror_sources": [
    {
      "resource_type": "LogicalPortMirrorSource",
      "port_ids": [
          "6a361832-43e4-430d-a48a-b84a6cba73c3"
      ]
    }
  ],
  "mirror_destination": {
    "resource_type": "LogicalPortMirrorDestination",
    "port_ids": [
        "3e42e8b2d-3047-4550-b230-dd1ee0e10b34"
    ]
  },
  "port_mirrorring_filters": [
      {
          "filter_action": "MIRROR",
          "src_ips": {
             "ip-addresses": [
                 "192.168.175.250",
                 "2001:bd6::c:2957:160:126"
             ]
```

```
            }
            "dst_ips": {
               "ip-addresses": [
                   "192.168.160.126",
                   "2001:bd6::c:2957:175:250"
               ]
            }
        }
    }
    "session_type": "LogicalPortMirrorSession",
    "preserve_original_vlan": false,
    "direction": "BIDIRECTIONAL",
    "_revision": 0
}
```

4 (Optional) You can call the `get mirroring-session <session-number>` CLI command to show the properties of the port mirroring session, including the filters.

Configure IPFIX in Manager Mode

IPFIX (Internet Protocol Flow Information Export) is a standard for the format and export of network flow information. You can configure IPFIX for switches and firewalls. For switches, network flow at VIFs (virtual interfaces) and pNICs (physical NICs) is exported. For firewalls, network flow that is managed by the distributed firewall component is exported.

NSX Cloud Note If using NSX Cloud, see NSX Features Supported with NSX Cloud for a list of auto-generated logical entities, supported features, and configurations required for NSX Cloud.

This feature is compliant with the standards specified in RFC 7011 and RFC 7012.

When you enable IPFIX, all configured host transport nodes will send IPFIX messages to the IPFIX collectors using port 4739. On ESXi hosts, NSX automatically opens port 4739.

IPFIX on ESXi sample tunnel packets in different ways. On ESXi the tunnel packet is sampled as two records:

- Outer packet record with some inner packet information
  - SrcAddr, DstAddr, SrcPort, DstPort, and Protocol refer to the outer packet.
  - Contains some enterprise entries to describe the inner packet.
- Inner packet record
  - SrcAddr, DstAddr, SrcPort, DstPort, and Protocol refer to the inner packet.

Configure Switch IPFIX Collectors in Manager Mode

You can configure IPFIX collectors for switches.

Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1. From your browser, log in with admin privileges to an NSX Manager at https://<nsx-manager-ip-address>.
2. Select **Plan & Troubleshoot > IPFIX**.
3. Click the **Switch IPFIX Collectors** tab.
4. Click **Add** to add a collector.
5. Enter a name and optionally a description.
6. Click **Add** and enter the IP address and port of a collector.

   You can add up to 4 collectors.
7. Click **Add**.

Configure Switch IPFIX Profiles in Manager Mode

You can configure IPFIX profiles for switches.

Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Select **Plan & Troubleshoot > IPFIX**.
3. Click the **Switch IPFIX Profiles** tab.
4. Click **Add** to add a profile.

Setting	Description
Name and Description	Enter a name and optionally a description. **Note** If you want to create a global profile, name the profile `Global`. A global profile cannot be edited or deleted from the UI, but you can do so using NSX APIs.
Active Timeout (seconds)	The length of time after which a flow will time out, even if more packets associated with the flow are received. Default is 300.
Idle Timeout (seconds)	The length of time after which a flow will time out, if no more packets associated with the flow are received. Default is 300.

Setting	Description
Export Overlay Flow	Setting that controls whether the sample result includes overlay flow information.
Sampling Probability (%)	The percentage of packets that will be sampled (approximately). Increasing this setting may have a performance impact on the hypervisors and collectors. If all hypervisors are sending more IPFIX packets to the collector, the collector may not be able to collect all packets. Setting the probability at the default value of 0.1% will keep the performance impact low.
Observation Domain ID	The observation domain ID identifies which observation domain the network flows originate from. Enter 0 to indicate no specific observation domain.
Collector Profile	Select a switch IPFIX collector that you configure in the previous step.
Priority	This parameter resolves conflicts when multiple profiles apply. The IPFIX exporter will use the profile with the highest priority only. A lower value means a higher priority.

5 Click **Add**.

Configure Firewall IPFIX Collectors in Manager Mode

You can configure IPFIX collectors for firewalls.

Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Plan & Troubleshoot > IPFIX**.

3 Click the **Firewall IPFIX Collectors** tab.

4 Click **Add** to add a collector.

5 Enter a name and optionally a description.

6 Click **Add** and enter the IP address and port of a collector.

You can add up to 4 collectors.

7 Click **Add**.

Configure Firewall IPFIX Profiles in Manager Mode

You can configure IPFIX profiles for firewalls.

Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Plan & Troubleshoot > IPFIX**.

3 Click the **Firewall IPFIX Profiles** tab.

4 Click **Add** to add a profile.

Setting	Description
Name and Description	Enter a name and optionally a description. **Note** If you want to create a global profile, name the profile `Global`. A global profile cannot be edited or deleted from the UI, but you can do so using NSX APIs.
Collector Configuration	Select a collector from the drop-down list.
Active Flow Export Timeout (Minutes)	The length of time after which a flow will time out, even if more packets associated with the flow are received. Default is 1.
Priority	This parameter resolves conflicts when multiple profiles apply. The IPFIX exporter will use the profile with the highest priority only. A lower value means a higher priority.
Observation Domain ID	This parameter identifies which observation domain the network flows originate from. The default is 0 and indicates no specific observation domain.

5 Click **Add**.

ESXi IPFIX Templates

An ESXi host transport node supports eight logical switch IPFIX flow templates and two distributed firewall IPFIX flow templates.

The following table lists VMware-specific elements in logical switch IPFIX packets.

Element ID	Parameter Name	Data Type	Unit
880	tenantProtocol	unsigned8	1 byte
881	tenantSourceIPv4	ipv4Address	4 bytes
882	tenantDestIPv4	ipv4Address	4 bytes
883	tenantSourceIPv6	ipv6Address	16 bytes
884	tenantDestIPv6	ipv6Address	16 bytes
886	tenantSourcePort	unsigned16	2 bytes

Element ID	Parameter Name	Data Type	Unit
887	tenantDestPort	unsigned16	2 bytes
888	egressInterfaceAttr	unsigned16	2 bytes
889	vxlanExportRole	unsigned8	1 byte
890	ingressInterfaceAttr	unsigned16	2 bytes
898	virtualObsID	string	variable length

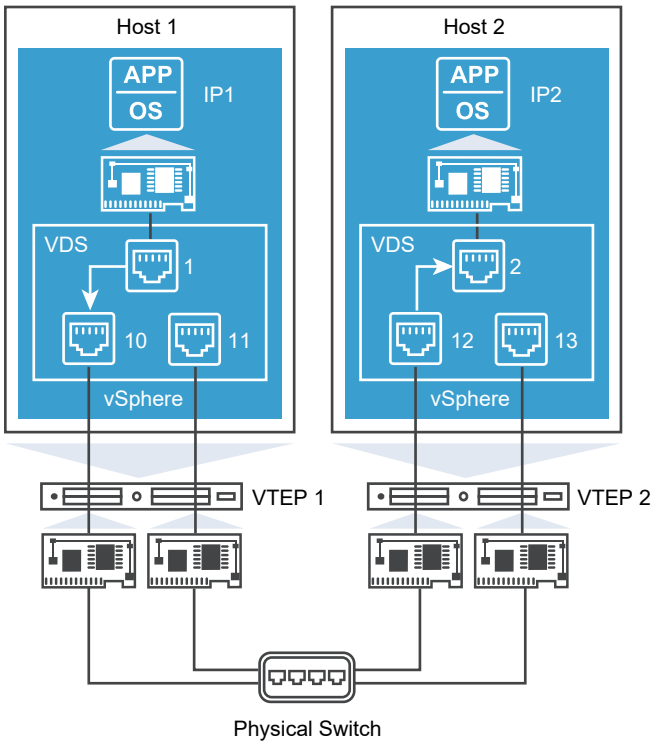
The following table lists VMware-specific elements in distributed firewall IPFIX packets.

Element ID	Parameter Name	Data Type	Unit
950	ruleId	unsigned32	4 bytes
951	vmUuid	string	16 bytes
952	vnicIndex	unsigned32	4 bytes
953	sessionFlags	unsigned8	1 byte
954	flowDirection	unsigned8	1 byte
955	algControlFlowId	unsigned64	8 bytes
956	algType	unsigned8	1 byte
957	algFlowType	unsigned8	1 byte
958	averageLatency	unsigned32	4 bytes
959	retransmissionCount	unsigned32	4 bytes
960	vifUuid	octetArray	16 bytes
961	vifId	string	variable length

ESXi Logical Switch IPFIX Templates

An ESXi host transport node supports eight logical switch IPFIX flow templates.

The following diagram shows the flow of traffic between VMs attached to ESXi hosts monitored by the IPFIX feature:



The IPv4 Encapsulated template will have the following elements:

- standard elements
- SrcAddr: VTEP1
- DstAddr: VTEP2
- tenantSourceIPv4: IP1
- tenantDestIPv4: IP2
- tenantSourcePort: 10000
- tenantDestPort: 80
- tenantProtocol: TCP
- ingressInterfaceAttr: 0x03 (tunnel port)
- egressInterfaceAttr: 0x01
- encapExportRole: 01
- virtualObsID: 89fd5032-2dc9-4fc3-993a-9bb4b616de54 (logical port ID)

IPv4 Template

Template ID: 256

```
IPFIX_TEMPLATE_START(IPFIX_FLOW_TYPE_IPv4)
IPFIX_TEMPLATE_FIELD(sourceIPv4Address, 4)
IPFIX_TEMPLATE_FIELD(destinationIPv4Address, 4)
IPFIX_TEMPLATE_FIELD(octetDeltaCount, 8)
```

```
IPFIX_TEMPLATE_FIELD(packetDeltaCount, 8)
IPFIX_TEMPLATE_FIELD(flowStartSysUpTime, 8)
IPFIX_TEMPLATE_FIELD(flowEndSysUpTime, 8)
IPFIX_TEMPLATE_FIELD(sourceTransportPort, 2)
IPFIX_TEMPLATE_FIELD(destinationTransportPort, 2)
IPFIX_TEMPLATE_FIELD(ingressInterface, 4)
IPFIX_TEMPLATE_FIELD(egressInterface, 4)
IPFIX_TEMPLATE_FIELD(encapId, 8)
IPFIX_TEMPLATE_FIELD(protocolIdentifier, 1)
IPFIX_TEMPLATE_FIELD(flowEndReason, 1)
IPFIX_TEMPLATE_FIELD(tcpFlags, 1)
IPFIX_TEMPLATE_FIELD(IPv4TOS, 1)
IPFIX_TEMPLATE_FIELD(maxTTL, 1)
IPFIX_TEMPLATE_FIELD(flowDir, 1)
// Specify the Interface port- Uplink Port, Access port,N.A
IPFIX_VMW_TEMPLATE_FIELD(ingressInterfaceAttr, 2)
IPFIX_VMW_TEMPLATE_FIELD(egressInterfaceAttr, 2)
IPFIX_VMW_TEMPLATE_FIELD(encapExportRole, 1)
IPFIX_VMW_TEMPLATE_VAR_LEN_FIELD(virtualObsID, virtualObsDataLen)
IPFIX_TEMPLATE_PADDING(paddingOctets, 1)
IPFIX_TEMPLATE_END()
```

IPv4 Encapsulated Template

Template ID: 257

```
IPFIX_TEMPLATE_START(IPFIX_FLOW_TYPE_IPv4_ENCAP)
IPFIX_TEMPLATE_FIELD(sourceIPv4Address, 4)
IPFIX_TEMPLATE_FIELD(destinationIPv4Address, 4)
IPFIX_TEMPLATE_FIELD(octetDeltaCount, 8)
IPFIX_TEMPLATE_FIELD(packetDeltaCount, 8)
IPFIX_TEMPLATE_FIELD(flowStartSysUpTime, 8)
IPFIX_TEMPLATE_FIELD(flowEndSysUpTime, 8)
IPFIX_TEMPLATE_FIELD(sourceTransportPort, 2)
IPFIX_TEMPLATE_FIELD(destinationTransportPort, 2)
IPFIX_TEMPLATE_FIELD(ingressInterface, 4)
IPFIX_TEMPLATE_FIELD(egressInterface, 4)
IPFIX_TEMPLATE_FIELD(protocolIdentifier, 1)
IPFIX_TEMPLATE_FIELD(flowEndReason, 1)
IPFIX_TEMPLATE_FIELD(tcpFlags, 1)
IPFIX_TEMPLATE_FIELD(IPv4TOS, 1)
IPFIX_TEMPLATE_FIELD(maxTTL, 1)
IPFIX_TEMPLATE_FIELD(flowDir, 1)
IPFIX_TEMPLATE_FIELD(encapId, 8)
IPFIX_VMW_TEMPLATE_FIELD(tenantSourceIPv4, 4)
IPFIX_VMW_TEMPLATE_FIELD(tenantDestIPv4, 4)
IPFIX_VMW_TEMPLATE_FIELD(tenantSourcePort, 2)
IPFIX_VMW_TEMPLATE_FIELD(tenantDestPort, 2)
IPFIX_VMW_TEMPLATE_FIELD(tenantProtocol, 1)
// Specify the Interface port - Uplink Port, Access port, N.A
IPFIX_VMW_TEMPLATE_FIELD(ingressInterfaceAttr, 2)
IPFIX_VMW_TEMPLATE_FIELD(egressInterfaceAttr, 2)
```

```
// TUNNEL-GW or no.
IPFIX_VMW_TEMPLATE_FIELD(encapExportRole, 1)
IPFIX_VMW_TEMPLATE_VAR_LEN_FIELD(virtualObsID, virtualObsDataLen)
IPFIX_TEMPLATE_END()
```

IPv4 ICMP Template

Template ID: 258

```
IPFIX_TEMPLATE_START(IPFIX_FLOW_TYPE_IPv4_ICMP)
IPFIX_TEMPLATE_FIELD(sourceIPv4Address, 4)
IPFIX_TEMPLATE_FIELD(destinationIPv4Address, 4)
IPFIX_TEMPLATE_FIELD(octetDeltaCount, 8)
IPFIX_TEMPLATE_FIELD(packetDeltaCount, 8)
IPFIX_TEMPLATE_FIELD(flowStartSysUpTime, 8)
IPFIX_TEMPLATE_FIELD(flowEndSysUpTime, 8)
IPFIX_TEMPLATE_FIELD(ingressInterface, 4)
IPFIX_TEMPLATE_FIELD(egressInterface, 4)
IPFIX_TEMPLATE_FIELD(protocolIdentifier, 1)
IPFIX_TEMPLATE_FIELD(flowEndReason, 1)
IPFIX_TEMPLATE_FIELD(IPv4TOS, 1)
IPFIX_TEMPLATE_FIELD(maxTTL, 1)
IPFIX_TEMPLATE_FIELD(flowDir, 1)
IPFIX_TEMPLATE_FIELD(encapId, 8)
// Specify the Interface port - Uplink Port, Access Port, or NA.
IPFIX_VMW_TEMPLATE_FIELD(ingressInterfaceAttr, 2)
IPFIX_VMW_TEMPLATE_FIELD(egressInterfaceAttr, 2)
IPFIX_VMW_TEMPLATE_FIELD(encapExportRole, 1)
IPFIX_VMW_TEMPLATE_VAR_LEN_FIELD(virtualObsID, virtualObsDataLen)
IPFIX_TEMPLATE_PADDING(paddingOctets, 2)
IPFIX_TEMPLATE_END()
```

IPv4 ICMP Encapsulated Template

Template ID: 259

```
IPFIX_TEMPLATE_START(IPFIX_FLOW_TYPE_IPv4_ICMP_ENCAP)
IPFIX_TEMPLATE_FIELD(sourceIPv4Address, 4)
IPFIX_TEMPLATE_FIELD(destinationIPv4Address, 4)
IPFIX_TEMPLATE_FIELD(octetDeltaCount, 8)
IPFIX_TEMPLATE_FIELD(packetDeltaCount, 8)
IPFIX_TEMPLATE_FIELD(flowStartSysUpTime, 8)
IPFIX_TEMPLATE_FIELD(flowEndSysUpTime, 8)
IPFIX_TEMPLATE_FIELD(sourceTransportPort, 2)
IPFIX_TEMPLATE_FIELD(destinationTransportPort, 2)
IPFIX_TEMPLATE_FIELD(ingressInterface, 4)
IPFIX_TEMPLATE_FIELD(egressInterface, 4)
IPFIX_TEMPLATE_FIELD(protocolIdentifier, 1)
IPFIX_TEMPLATE_FIELD(flowEndReason, 1)
IPFIX_TEMPLATE_FIELD(IPv4TOS, 1)
IPFIX_TEMPLATE_FIELD(maxTTL, 1)
IPFIX_TEMPLATE_FIELD(flowDir, 1)
IPFIX_TEMPLATE_FIELD(encapId, 8)
IPFIX_VMW_TEMPLATE_FIELD(tenantSourceIPv4, 4)
IPFIX_VMW_TEMPLATE_FIELD(tenantDestIPv4, 4)
IPFIX_VMW_TEMPLATE_FIELD(tenantProtocol, 1)
```

```
// Specify the Interface port- Uplink Port, Access port,N.A
IPFIX_VMW_TEMPLATE_FIELD(ingressInterfaceAttr, 2)
IPFIX_VMW_TEMPLATE_FIELD(egressInterfaceAttr, 2)
// TUNNEL-GW or no.
IPFIX_VMW_TEMPLATE_FIELD(encapExportRole, 1)
IPFIX_VMW_TEMPLATE_VAR_LEN_FIELD(virtualObsID, virtualObsDataLen)
IPFIX_TEMPLATE_PADDING(paddingOctets, 1)
IPFIX_TEMPLATE_END()
```

IPv6 Template

Template ID: 260

```
IPFIX_TEMPLATE_START(IPFIX_FLOW_TYPE_IPv6)
IPFIX_TEMPLATE_FIELD(sourceIPv6Address, 16)
IPFIX_TEMPLATE_FIELD(destinationIPv6Address, 16)
IPFIX_TEMPLATE_FIELD(octetDeltaCount, 8)
IPFIX_TEMPLATE_FIELD(packetDeltaCount, 8)
IPFIX_TEMPLATE_FIELD(flowStartSysUpTime, 8)
IPFIX_TEMPLATE_FIELD(flowEndSysUpTime, 8)
IPFIX_TEMPLATE_FIELD(sourceTransportPort, 2)
IPFIX_TEMPLATE_FIELD(destinationTransportPort, 2)
IPFIX_TEMPLATE_FIELD(ingressInterface, 4)
IPFIX_TEMPLATE_FIELD(egressInterface, 4)
IPFIX_TEMPLATE_FIELD(encapId, 8)
IPFIX_TEMPLATE_FIELD(protocolIdentifier, 1)
IPFIX_TEMPLATE_FIELD(flowEndReason, 1)
IPFIX_TEMPLATE_FIELD(tcpFlags, 1)
IPFIX_TEMPLATE_FIELD(IPv6TOS,1)
IPFIX_TEMPLATE_FIELD(maxTTL, 1)
IPFIX_TEMPLATE_FIELD(flowDir, 1)
// Specify the Interface port - Uplink Port, Access Port, or NA.
IPFIX_VMW_TEMPLATE_FIELD(ingressInterfaceAttr, 2)
IPFIX_VMW_TEMPLATE_FIELD(egressInterfaceAttr, 2)
IPFIX_VMW_TEMPLATE_FIELD(encapExportRole, 1)
IPFIX_VMW_TEMPLATE_VAR_LEN_FIELD(virtualObsID, virtualObsDataLen)
IPFIX_TEMPLATE_PADDING(paddingOctets, 1)
IPFIX_TEMPLATE_END()
```

IPv6 Encapsulated Template

Template ID: 261

```
IPFIX_TEMPLATE_START(IPFIX_FLOW_TYPE_IPv6_ENCAP)
IPFIX_TEMPLATE_FIELD(sourceIPv4Address, 4)
IPFIX_TEMPLATE_FIELD(destinationIPv4Address, 4)
IPFIX_TEMPLATE_FIELD(octetDeltaCount, 8)
IPFIX_TEMPLATE_FIELD(packetDeltaCount, 8)
IPFIX_TEMPLATE_FIELD(flowStartSysUpTime, 8)
IPFIX_TEMPLATE_FIELD(flowEndSysUpTime, 8)
IPFIX_TEMPLATE_FIELD(sourceTransportPort, 2)
IPFIX_TEMPLATE_FIELD(destinationTransportPort, 2)
IPFIX_TEMPLATE_FIELD(ingressInterface, 4)
IPFIX_TEMPLATE_FIELD(egressInterface, 4)
IPFIX_TEMPLATE_FIELD(protocolIdentifier, 1)
IPFIX_TEMPLATE_FIELD(flowEndReason, 1)
```

```
IPFIX_TEMPLATE_FIELD(tcpFlags, 1)
IPFIX_TEMPLATE_FIELD(IPv6TOS, 1)
IPFIX_TEMPLATE_FIELD(maxTTL, 1)
IPFIX_TEMPLATE_FIELD(flowDir, 1)
//ENCAP specific
IPFIX_TEMPLATE_FIELD(encapId, 8)
IPFIX_VMW_TEMPLATE_FIELD(tenantSourceIPv6, 16)
IPFIX_VMW_TEMPLATE_FIELD(tenantDestIPv6, 16)
IPFIX_VMW_TEMPLATE_FIELD(tenantSourcePort, 2)
IPFIX_VMW_TEMPLATE_FIELD(tenantDestPort, 2)
IPFIX_VMW_TEMPLATE_FIELD(tenantProtocol, 1)
// Specify the Interface port - Uplink Port, Access Port, or NA
IPFIX_VMW_TEMPLATE_FIELD(ingressInterfaceAttr, 2)
IPFIX_VMW_TEMPLATE_FIELD(egressInterfaceAttr, 2)
// TUNNEL-GW or no.
IPFIX_VMW_TEMPLATE_FIELD(encapExportRole, 1)
IPFIX_VMW_TEMPLATE_VAR_LEN_FIELD(virtualObsID, virtualObsDataLen)
IPFIX_TEMPLATE_END()
```

IPv6 ICMP Template

Template ID: 262

```
IPFIX_TEMPLATE_START(IPFIX_FLOW_TYPE_IPv6_ICMP)
IPFIX_TEMPLATE_FIELD(sourceIPv6Address, 16)
IPFIX_TEMPLATE_FIELD(destinationIPv6Address, 16)
IPFIX_TEMPLATE_FIELD(octetDeltaCount, 8)
IPFIX_TEMPLATE_FIELD(packetDeltaCount, 8)
IPFIX_TEMPLATE_FIELD(flowStartSysUpTime, 8)
IPFIX_TEMPLATE_FIELD(flowEndSysUpTime, 8)
IPFIX_TEMPLATE_FIELD(ingressInterface, 4)
IPFIX_TEMPLATE_FIELD(egressInterface, 4)
IPFIX_TEMPLATE_FIELD(protocolIdentifier, 1)
IPFIX_TEMPLATE_FIELD(flowEndReason, 1)
IPFIX_TEMPLATE_FIELD(IPv6TOS, 1)
IPFIX_TEMPLATE_FIELD(maxTTL, 1)
IPFIX_TEMPLATE_FIELD(flowDir, 1)
IPFIX_TEMPLATE_FIELD(encapId, 8)
// Specify the Interface port - Uplink Port, Access Port, or NA.
IPFIX_VMW_TEMPLATE_FIELD(ingressInterfaceAttr, 2)
IPFIX_VMW_TEMPLATE_FIELD(egressInterfaceAttr, 2)
IPFIX_VMW_TEMPLATE_FIELD(encapExportRole, 1)
IPFIX_VMW_TEMPLATE_VAR_LEN_FIELD(virtualObsID, virtualObsDataLen)
IPFIX_TEMPLATE_PADDING(paddingOctets, 2)
IPFIX_TEMPLATE_END()
```

IPv6 ICMP Encapsulated Template

Template ID: 263

```
IPFIX_TEMPLATE_START(IPFIX_FLOW_TYPE_IPv6_ICMP_ENCAP)
IPFIX_TEMPLATE_FIELD(sourceIPv4Address, 4)
IPFIX_TEMPLATE_FIELD(destinationIPv4Address, 4)
IPFIX_TEMPLATE_FIELD(octetDeltaCount, 8)
IPFIX_TEMPLATE_FIELD(packetDeltaCount, 8)
IPFIX_TEMPLATE_FIELD(flowStartSysUpTime, 8)
```

```
IPFIX_TEMPLATE_FIELD(flowEndSysUpTime, 8)
IPFIX_VMW_TEMPLATE_FIELD(sourceTransportPort, 2)
IPFIX_VMW_TEMPLATE_FIELD(destinationTransportPort, 2)
IPFIX_TEMPLATE_FIELD(ingressInterface, 4)
IPFIX_TEMPLATE_FIELD(egressInterface, 4)
IPFIX_TEMPLATE_FIELD(protocolIdentifier, 1)
IPFIX_TEMPLATE_FIELD(IPv6TOS, 1)
IPFIX_TEMPLATE_FIELD(maxTTL, 1)
IPFIX_TEMPLATE_FIELD(flowDir, 1)
IPFIX_TEMPLATE_FIELD(flowEndReason, 1)
//ENCAP Specific
IPFIX_TEMPLATE_FIELD(encapId, 8)
IPFIX_VMW_TEMPLATE_FIELD(tenantSourceIPv6, 16)
IPFIX_VMW_TEMPLATE_FIELD(tenantDestIPv6, 16)
IPFIX_VMW_TEMPLATE_FIELD(tenantProtocol, 1)
// Specify the Interface port - Uplink Port, Access Port, or NA
IPFIX_VMW_TEMPLATE_FIELD(ingressInterfaceAttr, 2)
IPFIX_VMW_TEMPLATE_FIELD(egressInterfaceAttr, 2)
// TUNNEL-GW or no.
IPFIX_VMW_TEMPLATE_FIELD(encapExportRole, 1)
IPFIX_VMW_TEMPLATE_VAR_LEN_FIELD(virtualObsID, virtualObsDataLen)
IPFIX_TEMPLATE_PADDING(paddingOctets, 1)
IPFIX_TEMPLATE_END()
```

ESXi Distributed Firewall IPFIX Templates

An ESXi host transport node supports two distributed firewall IPFIX flow templates.

IPv4 Template

Template ID: 288

```
IPFIX_TEMPLATE_FIELD(sourceIPv4Address,4)
IPFIX_TEMPLATE_FIELD(destinationIPv4Address,4)
IPFIX_TEMPLATE_FIELD(sourceTransportPort,2)
IPFIX_TEMPLATE_FIELD(destinationTransportPort,2)
IPFIX_TEMPLATE_FIELD(protocolIdentifier,1)
IPFIX_TEMPLATE_FIELD(icmpTypeIPv4,1)
IPFIX_TEMPLATE_FIELD(icmpCodeIPv4,1)
IPFIX_TEMPLATE_FIELD(flowStartSeconds,4)
IPFIX_TEMPLATE_FIELD(flowEndSeconds,4)
IPFIX_TEMPLATE_FIELD(octetDeltaCount,8)
IPFIX_TEMPLATE_FIELD(packetDeltaCount,8)
IPFIX_TEMPLATE_FIELD(firewallEvent,1)
IPFIX_TEMPLATE_FIELD(direction,1)
IPFIX_TEMPLATE_FIELD(ruleId,4)
IPFIX_TEMPLATE_FIELD(vifUuid,16)
IPFIX_TEMPLATE_FIELD(sessionFlags,1)
IPFIX_TEMPLATE_FIELD(flowDirection,1)
IPFIX_TEMPLATE_FIELD(flowId,8)
IPFIX_TEMPLATE_FIELD(algControlFlowId,8)
IPFIX_TEMPLATE_FIELD(algType,1)
IPFIX_TEMPLATE_FIELD(algFlowType,1)
IPFIX_TEMPLATE_FIELD(averageLatency,4)
IPFIX_TEMPLATE_FIELD(retransmissionCount,4)
```

IPv6 Template

Template ID: 289

```
IPFIX_TEMPLATE_FIELD(sourceIPv6Address,16)
IPFIX_TEMPLATE_FIELD(destinationIPv6Address,16)
IPFIX_TEMPLATE_FIELD(sourceTransportPort,2)
IPFIX_TEMPLATE_FIELD(destinationTransportPort,2)
IPFIX_TEMPLATE_FIELD(protocolIdentifier,1)
IPFIX_TEMPLATE_FIELD(icmpTypeIPv6,1)
IPFIX_TEMPLATE_FIELD(icmpCodeIPv6,1)
IPFIX_TEMPLATE_FIELD(flowStartSeconds,4)
IPFIX_TEMPLATE_FIELD(flowEndSeconds,4)
IPFIX_TEMPLATE_FIELD(octetDeltaCount,8)
IPFIX_TEMPLATE_FIELD(packetDeltaCount,8)
IPFIX_TEMPLATE_FIELD(firewallEvent,1)
IPFIX_TEMPLATE_FIELD(direction,1)
IPFIX_TEMPLATE_FIELD(ruleId,4)
IPFIX_TEMPLATE_FIELD(vifUuid,16)
IPFIX_TEMPLATE_FIELD(sessionFlags,1)
IPFIX_TEMPLATE_FIELD(flowDirection,1)
IPFIX_TEMPLATE_FIELD(flowId,8)
IPFIX_TEMPLATE_FIELD(algControlFlowId,8)
IPFIX_TEMPLATE_FIELD(algType,1)
IPFIX_TEMPLATE_FIELD(algFlowType,1)
IPFIX_TEMPLATE_FIELD(averageLatency,4)
IPFIX_TEMPLATE_FIELD(retransmissionCount,4)
```

Monitor a Logical Switch Port Activity in Manager Mode

You can monitor the logical port activity for example, to troubleshoot network congestion and packets being dropped

Prerequisites

- Verify that a logical switch port is configured. See Connecting a VM to a Logical Switch in Manager Mode.
- Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Verify that a logical switch port is configured. See Connecting a VM to a Logical Switch in Manager Mode.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Select **Networking > Logical Switches > Ports**
3. Click the name of a port.

4 Click the **Monitor** tab.

   The port status and statistics are displayed.

5 To download a CSV file of the MAC addresses that has been learned by the host, click **Download MAC Table**.

6 To monitor activity on the port, click **Begin Tracking**.

   A port tracking page opens. You can view the bidirectional port traffic and identify dropped packets. The port tracker page also lists the switching profiles attached to the logical switch port.

Results

If you notice dropped packets because of network congestion, you can configure a QoS switching profile for the logical switch port to prevent data loss on preferred packets. See Understanding QoS Switching Profile.

Checking CPU Usage and Network Latency

Use APIs to monitor CPU usage and network latency.

Checking CPU Usage

These APIs work similar to the `top` command:

- Run the following API to check process status for all appliances:

```
GET /api/v1/systemhealth/appliances/process/status
```

- Run the following API to check the process status of a specific appliance.

```
GET https://<nsx-manager-ip>/api/vi/systemhealth/appliances/<appliance-id>/process/status
```

Checking Network Latency

- Run the following API to check for network latency values for all appliances:

```
GET /api/v1/systemhealth/appliances/latency/status
```

- Run the following API to check for network latency values for a specific appliance:

```
GET https://<nsx-manager-ip>/api/v1/systemhealth/appliances/<appliance-id>/latency/status
```

See the latest version of the *NSX API Guide* at https://code.vmware.com/ for API details.

Live Traffic Analysis

Live Traffic Analysis (LTA) provides helpful insight about tracing live traffic and bi-directional packet tracing. Traffic analysis monitors live traffic at a source or between source and destination along with the packet capture. You can identify bad flows between the source and the destination. If the packet counter of the certain flow at the source endpoint is much higher than the packet counter of the certain flow at the destination endpoint, packet drop may occur between two endpoints. Hence the flow is probably a bad flow which you can trace for further analysis. Thus traffic analysis is helpful in troubleshooting virtual network issues. You can find the number of packet enter or leave a port, and the unexpected packet drop.

Limitations

- LTA supports only on the overlay-backed NSX environments.
- LTA is not supported on DPU.
- LTA is not supported on T0 Active/Active setup.
- LTA is not supported on Global Manager.
- LTA cannot observe VMC components that do not belong to the NSX management domain, such as IGW.

Create a Session

To start a LTA session, click **New session**. For details, see Create a Live Traffic Analysis Session.

Session List

After you create a session, you can view the list of all the active sessions. The session persists only for one hour. The session status can be as follows:

Status	Description
In Progress	The session is collecting the results. Wait for the session to finish.
Realized	Realized session is that session whose intent configuration is realized on the management plane.
Unrealized	Unrealized session is that session whose intent configuration is yet to be realized on management plane.
Partial Finished	Only partial session result is available. Some session results might have lost.
Finished	After the session is finished, you can view the session for further analysis. You can also rerun, duplicate, or delete the session.
Cancelled	The session is canceled by exception.

Status	Description
Invalidated	The session is canceled proactively by the data plane service because of disconnection of the source port or the data plane service is down. Make sure that the source or destination port is not migrated or disconnected during the running LTA session and retry.
Timeout	The session is timed out.

To view the latest status, click the **Refresh** icon. To perform the following tasks, click  and then click the required option.

Option	Description
Rerun	Run an existing session again. The session persists only for one hour.
Delete	Delete an existing session.
Copy Path to Clipboard	Get the path of the LTA configuration. You can use the copied path later. You can filter the LTA sessions using the copied path.

Session Details

After the session is finished, you can click the session ID link and perform the traffic analysis. You can view data under the **Observations** tab and the **Packet Count** tab.

Observations Tab	Description
Delivered	Total number of received observations for the traceflow round.
Dropped	Total number of dropped observations for the traceflow round.
Physical Hop Count	The sequence number is the traceflow observation hop count. The hop count for observations on the transport node that a traceflow packet is injected in can be 0. The hop count is incremented each time a subsequent transport node receives the traceflow packet. The sequence number of 999 indicates that the hop count could not be determined for the containing observation.
Observation Type	The observation type can be Forwarded, Dropped, Delivered, Received, and Injected.
Transport Node	The name of the transport node that observed a traceflow packet.
Component	The name of the component that issued the observation.
Actions	You can view details for certain traceflow observation like MAC details of a logical switch or segment.

Packet Count Tab	Description
Component	The type of the component.
Transport Node	ID of the transport node that observed a traceflow packet.
Packets Received, Forwarded, and Dropped	Number of traceflow packets that were received, forwarded, or dropped.

Create a Live Traffic Analysis Session

You can create a Live Traffic Analysis (LTA) session. Traffic analysis is helpful in troubleshooting virtual network problems.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Navigate to **Plan & Troubleshoot > Traffic Analysis > Live Traffic Analysis > Get Started**.

3 To start a traffic analysis, click **New Session**.

4 System generates the session name. If you wish, you can edit the session name.

5 Select the trace, packet count, and packet capture and provide related information.

Protocol	Parameters
Trace & Packet Capture Sampling Type	System supports only the *FirstNSampling* type. First *N* number of packets that match the packet filter under the **Advanced Settings** tab are sampled for analysis.
Trace	By default, the toggle is active. Trace generates the packet traceflow. If you select only the source, system generates the traceflow for the packets ingressed from the source (matching the forward filter, if any). If you also select the destination, in addition system generates the traceflow for the packets ingressed from the destination (matching the reverse filter, if any).
Trace Sampling Value	Number of packets to be sampled. Enter the value from 1 through 50.
Packet Count	Captures the counters at each of the observation points when packet is received. By default, the toggle is active.
Packet Capture	Generates the PCAP files with live trace telemetry. By default, the toggle is active.
Packet Capture Sampling Value	Enter the value from 1 through 500.

6 Specify the source and destination information according to the traffic type. Click **Add Destination**, and select a destination to capture.

Bi-directional trace traces the traffic ingressed from the source and the traffic ingressed from the destination, respectively.

For example, consider three VMs: *App-VM*, *Web-VM*, and *DB-VM*, and there is a ping traffic from the *App-VM* to the *Web-VM* and a ping traffic from the *App-VM* to the *DB-VM*. If you perform bi-directional trace without any packet filters with source as the *App-VM* and destination as the *Web-VM*, then the traces of the forward direction contains traces for ICMP echo from the *App-VM* to the *Web-VM* and the *DB-VM*, while the traces of the reverse direction contains traces for ICMP reply from the *Web-VM* to the *App-VM*.

To see the traffic between the source and destination only, specify the proper packet filter.

Traffic Type	Source	Destination
Virtual Machine	For a VM: ■ Select the name of the virtual machine from the list. ■ You can select or view the virtual interface and the segment port for the selected VM.	For a VM: ■ Select the name of the virtual machine from the list. ■ You can select or view the virtual interface and the segment port for the selected VM.
Port/ Interface	For a logical port: ■ Select an attachment type as **Virtual Interface**, **Edge Uplink**, **Edge Centralized Service**, or **IPSec**. ■ Select a port. **IPSec** session lists only the route-based VPN sessions.	For a logical port: ■ Select an attachment type as **Virtual Interface**, **Edge Uplink**, **Edge Centralized Service**, or **IPSec**. ■ Select a port.

Note

- You cannot configure a source attachment type as **Edge Uplink** and a destination attachment as **IPSec**, or vice-a-versa.
- When you add **IPSec**, **Edge Uplink**, or **VM** as a destination attachment type, you must configure a source attachment type.
- When you add **IPSec** attachment as an intermediate interface, you might see observations for both destinations- inbound as well as outbound traffic.

7 (Optional) Click **Advanced Settings** and view the advanced options.

Enter the desired values for the following parameters and click **Apply**.

Table 21-1.

Option	Description
General Tab	
Session Timeout	Default timeout value is 10 seconds. You can add value between 5 to 300 seconds.
Packet Count Settings > Include Only Interface Points	Activate the toggle button if you want the packet count observation to include only the interface points.
Filters Tab	
IP Type	Select **IPv4** or **IPv6**.

Table 21-1. (continued)

Option	Description
Forward Filters or Reverse Filters	Forward filters formulate the flows of interest for the traffic ingressed from the source. Reverse filters formulate the flows of interest for the traffic ingressed from the destination. You can apply filter based on 5-tuple or plain text as **Fields Filter Data** or **Plain Filter Data**.
Filter Type: Fields Filter Data	For bi-directional LTA session, the system populates the source and destination IPs. Select the protocol type and enter the source and destination IPs and port details. If you change the filter type to plain filter data and set the empty values, then all IP observations are reported for all the traffic in the source and destination VMs. ■ If you select logical port with **Edge Uplink** attachment, then select filter type as **Fields Filter Data**. If your protocol type is **ESP**, then select the IPSec session name and enter the Service Path Index (SPI) value. SPI is an optional field. ■ If you select logical port with **IPSec** attachment, then select filter type as **Fields Filter Data** and enter the required values.
Filter Type: Plain Filter Data	Add details for basic and extended filter such as IP address and port number. **Note** Plain filter data is realized on ESXi only.

8 Click **Start Session**.

Results

You can view the status of the created session. After the session is finished, you can view the session for further analysis.

What to do next

You can perform the following tasks:

Option	Description
Download PCAP Files	You can download the PCAP file to your system for further analysis. For the bi-directional LTA sessions, you can download both forward and reverse PCAP files.
Rerun	Run an existing session again. The session persists only for one hour.
Duplicate Session	You can copy the session parameters to create a new session. You can quickly change few options in the new session.
New Trace	You can start a new traffic analysis session again.

Authentication and Authorization 22

You can log in to NSX Manager using a local user account, a user account managed by VMware Identity Manager (vIDM), or a user account managed by a directory service such as Active Directory over LDAP or OpenLDAP. You can also assign roles to user accounts managed by vIDM or a directory service to implement role-based access control.

NSX Manager recognizes only system-generated session identifiers and invalidates session identifiers upon administrator logout or other session termination. Upon successful login, the NSX Manager uses a random number generator to create a random session ID and stores that ID in memory. When clients make requests to the NSX Manager, it only allows clients to authenticate if the session ID they present matches one of the IDs generated by the server. When any user logs out of NSX Manager, the session identifier is immediately destroyed and cannot be reused.

Access to NSX Manager via UI, API and CLI is subject to authentication and authorization. In addition, such access will generate audit logs. This logging is enabled by default and cannot be disabled. Auditing of sessions is initiated at system startup. Audit log messages include the text `audit="true"` in the structured data part of the log message.

Local user passwords on NSX appliances are secured using the default Linux/PAM libraries which store the hashed and salted representation in `/etc/shadow`. NSX Manager uses the SHA512 cryptographic hash algorithm to hash the local user passwords. During authentication, the password entered by the user is obfuscated. Other passwords are encrypted using a random key that is stored in the local file system. For more details, see the VMware Security Hardening Guides or review the SHA512 Ubuntu MAN pages and the Internet FAQ titled "Understanding /etc/shadow file format on Linux."

This chapter includes the following topics:

- Managing Local User Accounts
- Password Management
- Authentication Policy Settings
- Integration with VMware Identity Manager/Workspace ONE Access
- Integration with LDAP
- Configuring Both vIDM and LDAP or Transitioning from vIDM to LDAP
- NSX API Authentication Using a Session Cookie

- Add a Role Assignment or Principal Identity
- Role-Based Access Control
- Create or Manage Custom Roles
- Logging User Account Changes

Managing Local User Accounts

Starting with NSX 4.1, in addition to the admin user in the Enterprise environment, any user with the Enterprise Admin role can add new local users and remove audit and guest users. To administer NSX Manager, you must log in as admin. Some account management privileges are extended to the Enterprise Admin role.

After installation, the admin account is active. To use other local user accounts, you must activate all other local user accounts including the audit account.

Each NSX appliance has four default local accounts including **admin**, **audit**, and two local guest user accounts. The two default local user accounts are **guestuser1** and **guestuser2**. In the NSX Cloud environment, the user accounts are **cloud_admin** and **cloud_audit**.

- The local guest user accounts, including the cloud_admin and the cloud_audit accounts, have the Auditor role. You can change their role assignments.
- You can change role assignments for the guest users. For example, you can mix read-only and full access permissions by creating a clone Auditor role.
- The admin or the account owners can reset local user account passwords.
- Starting with NSX 4.1, you can add up to 10 guest users with a maximum of 14 total users.
- You can delete the audit user and the guest users. You cannot delete the admin user.

An NSX appliance also has the **root** user account. You cannot log in to the NSX Manager UI as root, and you cannot manage this account through the NSX Manager UI. The root user can log in to an appliance through the CLI, but cannot use the NSX CLI commands. The root user account cannot be renamed, deactivated, or deleted.

The root user has special privileges. You must not log in to an NSX appliance as root and make changes that are not documented in this guide, except when under the guidance of VMware. Changes made by the root user can cause catastrophic failures. In a production environment, the root password must be secure and made available for privileged access only.

For details on how to manage your local user accounts, including password reset, adding, and deleting users, see Manage Local User Accounts . For additional security-related information about the NSX Manager, see the section "Security" in Chapter 1 NSX Manager.

Add a Local User

Starting in NSX 4.1, in addition to admin, users with the Enterprise Admin role can create new local user accounts in the NSX Manager.

You can complete this task using the UI, API or CLI. The Cloud admin users do not have privileges to add new local users.

Prerequisites

- Only a user with the Enterprise Admin role can create new local guest user accounts or replace a deleted audit user account.
- The maximum number of local guest users is 14.
- The maximum number of audit users is one. Guest users always have the default Auditor role.

Procedure

1. From your browser, log in to an NSX Manager at https://<nsx-manager-ip-address>.
2. Select **System > User Management > Local Users > Add**.
   a. To add an audit user, select **Audit User**.

      If an audit user exists, you must first delete that user and retry the step. Only one audit user is available.

   b. To add a guest user, select **Local User**.

      If the maximum number of guest users already exists, then the Add button is inactive. Delete one or more of the guest users and try again.
3. Click **Save**.
4. (Optional) To activate a new local user account, locate the user name and click  .
   a. Select **Activate User**.
   b. Enter a password for the user.
   c. Click **Save**.

   You can choose to activate your local user at a later time, but the user cannot log in until their status is active.

Results

1. The newly added user displays in the Local Users list with the default Auditor role.
2. If you activate the user, the new user must change their password after logging in. Refer them to Manage Local User Accounts for details.

What to do next

To read more about how to create a custom role or add more roles to a local user, go to step 4 in Manage Local User Accounts . If you plan to activate the user at a later time, refer to Activate a Local User in NSX Manager.

Activate a Local User in NSX Manager

You can activate a local user account from the NSX Manager UI. Initially, only the `admin` account is active. You can authorize additional user access by activating the other local user accounts. As part of the new user activation, you can also optionally change the username and roles.

Prerequisites

Familiarize yourself with the password complexity requirements for NSX Manager and NSX Edge. See "NSX Manager Installation" and "NSX Edge Installation" in the *NSX Installation Guide*. Guest users are available on NSX Manager only.

Procedure

1. From your browser, log in as admin to an NSX Manager at https://<nsx-managr-ip-address>.
2. Select **System > Users and Roles > Local Users**.
3. Locate the local user name you want to add and click  .
4. Select **Activate**. The guest user and audit accounts are inactive by default and must be activated through the UI and API before using.

   New expiration dates are created after user activation. To view or change password expiration as admin, see Manage Local User Accounts .
5. (Optional) To change the local user name:
   a. Click  for that user and select **Edit**.
   b. Enter the user name changes.
   c. Click **Save**.
6. To change the user role assigned to one or both of the guest users:
   a. On the **User Role Assignment** tab, click  and select **Edit**.
   b. Select the role change from the drop down list.
   c. Click **Save**.
7. (Optional) To deactivate any of the local users, click  for that user, select **Deactivate User**, and click **Deactivate**.

   For more details on customizing roles, see Create or Manage Custom Roles.

Manage Local User Accounts

You can manage the default and new local guest users, through the NSX Manager UI.

You cannot deactivate the admin user or change their role assignments. You also cannot change the audit user role assignments. The admin user or any user with the Enterprise Admin role can perform the following tasks unless noted otherwise:

- Starting with NSX 4.1, users with the Enterprise Admin role can add a user. To add users, go to Add a Local User .
- Starting with NSX 4.1, the admin user can delete a user. To delete or remove a user, go to Delete a Local User.
- Activate or deactivate any local user accounts, except for admin.
- Change user role assignments for the two guest users.
- Add a new role, clone an existing role, and edit or delete user-created roles.
- Reset user passwords. In addition, all local users can reset their own passwords if they know their current password.
- Change the user names for any of the local user accounts.
- Change password expiration settings.

The audit and guest users have default read privileges to the NSX environment and are not active by default. Before they can log in to NSX Manager, you must activate the accounts first.

Any changes to the local user accounts is audited.

By default, user passwords expire after 90 days. You can change or deactivate the password expiration for each user.

When a user logs in to NSX Manager, if the password is set to expire within 30 days, the NSX Manager UI displays a password expiration notification. If you set the password expiration to 30 days or less, the notification is always present. The notification includes a **Change Password** link. Click the link to change the user's password.

Prerequisites

Familiarize yourself with the password complexity requirements for NSX Manager and NSX Edge. See "NSX Manager Installation" and "NSX Edge Installation" in the *NSX Installation Guide*. To change password complexity, review Authentication Policy Settings

Procedure

1. From your browser, log in as admin to an NSX Manager at https://<nsx-manager-ip-address>.
2. Select **System > User Management**.
3. To activate a user, select the **Local Users** tab and locate the user name.
   a. Click  .
   b. Select **Activate User**.

c Enter a password for the user.

d Click **Save**.

4 To change or reset a user password, select the **Local Users** tab and locate the user name.

a Click  .

b Select **Reset Password**.

c Enter the password details.

d Click **Save**.

5 (Optional) To add a user role assignment for a guest user in addition to the existing Auditor role, select the **User Role Assignment** tab and locate the user name.

a Click  .

b Select **Edit**.

c Click the highlighted number link under Roles. This is the number of user roles assigned to this user.

d From the Set Roles/Scope window, click **Add Role**.

e Select the role from the dropdown list.

Depending on the role you select, you may need to click **Set** to select the scope for this role assignment. If you use projects and want to specify specific access, select the projects to apply to this user role.

f To add the new role, click **Add**, then click **Apply**.

g Click **Save**.

6 To edit a user role assignment for a guest user, select the **User Role Assignment** tab and locate the user name.

a Click  .

b Select **Edit**.

c Click the highlighted number link under Roles. This is the number of user roles assigned to this user.

From the Set Roles/Scope window, you can add a role, edit an existing role, or delete a role.

d Update this guest user roles, then click **Apply**.

e Click **Save**.

7 To delete a role assignment, select the **User Role Assignment** tab and locate the user name.

   a Click  .

   b Select **Edit**.

   c Click the highlighted link under Roles.

   This is the number of user roles of this user.

   d Click  .

   e Select **Delete** and confirm the deletion.

   f Click **Apply**.

   g Click **Save**.

8 (Optional) To change a user name, select the **Local Users** tab and locate the user name.

   a Click  .

   b Select **Edit**.

   c Change the user name.

   d Click **Save** and **Continue**.

   If you change the audit user name, you can identify it by the Auditor label containing a padlock.

9 (Optional) To deactivate a user, select the **Local Users** and locate the user name.

   a Click  .

   b Select **Deactivate User** and confirm.

   c Click **Deactivate**.

10 To get the password expiration information, from the **Local Users** tab, expand the row for the user that you want to view.

11 (Optional) To change the password expiration settings, log in to the appliance's CLI as **admin**.

   a To set the password expiration time in days, run the `set user <username> password-expiration <number of days>` command.

```
nsxcli> set user admin password-expiration 120
nsxcli>
```

   b To deactivate password expiration, run the `clear user <username> password-expiration`

```
nsxcli> clear user admin password-expiration
nsxcli>
```

Manage Local User's Password or Name Using the CLI

You can manage NSX Manager user accounts through an NSX appliance's CLI. This topic describes how the admin user manages user account details using the CLI. Alternately, you can also use the UI.

The `admin` user can manage passwords, change the name of the admin and other users, and add, delete, or deactivate users. Any user account change is audited.

For extended access, see Manage Local User Accounts .

The audit user has read privileges to the NSX environment and is not active by default unless the audit password is provided during NSX installation. To activate the audit user after installation, use the UI or log in as admin to the CLI and run the `set user audit password` command and provide a new password.

By default, user passwords expire after 90 days. You can change or deactivate the password expiration for each user.

When an NSX Manager user password is within 30 days of expiring, the NSX Manager UI displays a password notification after logging in. The notification includes a **Change Password** link. To change the password, click the link.

For CLI details, see the *NSX Command-Line Interface Reference*.

Prerequisites

Familiarize yourself with the password complexity requirements for NSX Manager and NSX Edge. See "NSX Manager Installation" and "NSX Edge Installation" in the *NSX Installation Guide*.

Procedure

1. Log in to the CLI of the appliance as `admin`.
2. To change a user password, run the `set user <username> password` command. For example:

```
nsx> set user audit password
New password:
Confirm new password:
nsx>
```

3. To change the name of a user, run the `set user <username> username <new username>` command. For example:

```
nsx> set user admin username admin1
nsx>
```

4 To see a list of existing user names, run the `set user [TAB][TAB]` command. For example:

```
nsx> set user [TAB][TAB]
  admin     Username of user
  audit     Username of user
  root      Username of user
```

5 To get password expiration information, run the `get user <username> password-expiration` command. For example:

```
nsx> get user admin password-expiration
```

```
Tue Jun 07 2022 UTC 06:33:29.963
Password expires 90 days after last change
 Current password will expire in 90 days
User will receive warning messages 7 days before password expires.
Password expires 90 days after last change
nsx>
```

6 To set the password expiration time in days, run the `set user <username> password-expiration <number of days>` command. For example:

```
nsx> set user admin password-expiration 120
nsx>
```

7 To deactivate password expiration for users, run the `clear user <username> password-expiration` command. For example:

```
nsx> clear user admin password-expiration
nsx>
```

8 To change the default number of days a user receives a warning message prior to their password expiration, run the `set user <username> password-expiration-warning <password-expiration-warn-days>` command. Default is 7. Range is 1 to 1999. For example:

```
set user admin password-expiration-warning 14
```

Delete a Local User

You can remove the local user accounts from NSX. These include the audit and guest user accounts. You can recreate deleted local user accounts as needed.

You can complete this task using the UI, API or CLI. Other users, including the Enterprise or Cloud admin users, do not have privileges to delete local users.

Prerequisites

- Only a user with the admin role can delete local user accounts or replace deleted audit user and guest user accounts.

- You can remove Edge node local users, but you cannot add them back on the Edge.

Procedure

1. From your browser, log in as admin to an NSX Manager at https://<nsx-manager-ip-address>.
2. Select **System > User Management > Local Users**.
3. To remove a local user account, locate the user name you want to remove and click  .
4. From the popup menu, click **Delete**.
5. To confirm, click **Delete** again.

   A message displays that your deletion was successful.

Password Management

This topic provides links to the many options to manage your user passwords using the NSX user interface, CLI, and API.

Resetting User Passwords

There are many ways to reset your user password or have an administrator reset it for you. This topic describes the scenarios and links to the procedures to complete your password resets.

Prerequisites

To resolve your password issues, use the following table to find the link to the appropriate procedure.

These password rules are followed in NSX:

- Password changes in the UI may require the user's current password.
- When the admin resets an audit or a local user's password, it results in the immediate expiration of the admin-generated password that forces the user to change the password during login.
- Expired passwords are reset by providing the existing expired password.

User	Scenario	UI	CLI	API
You are the Admin or the Enterprise Admin	Reset your own forgotten password		If you know the root password, use the NSX CLI. See Resetting the Passwords of an Appliance.	
	Reset any user password (logged in as Admin or Enterprise Admin)	See Manage Local User Accounts . To set a user password during activation, see Activate a Local User in NSX Manager.	Use the NSX CLI (except guest users). See Manage Local User's Password or Name Using the CLI Admin only; remote Enterprise Admin does not have access.	Go to https://code.vmware.com/apis/1163/nsx-t and select your release number. Search under **System Administration > Configuration > Fabric > Nodes > User Management > Users** for API details.
	Reset expired password (Admin only access)	To reset other user's expired passwords, see Resetting Expired Passwords. Admin cannot reset their own expired password in the UI.	■ To reset admin expired password, SSH or log in to the NSX CLI. Requires existing password. ■ Only admin can reset the Enterprise admin's expired password. ■ To reset the audit user expired password, SSH or log in to the NSX CLI. See Manage Local User's Password or Name Using the CLI.	Admin can go to https://code.vmware.com/apis/1163/nsx-t and select your release number. Search for **System Administration > Configuration > Fabric > Nodes > User Management** to find the API details.
You are the Audit user	Reset your own forgotten password	Contact your administrator.	Contact your administrator.	Contact your administrator.
	Reset your own password (requires known password)	If you know your password, see Manage Local User Accounts .	Contact your administrator.	Go to https://code.vmware.com/apis/1163/nsx-t and select your release number. Search for **System Administration > Configuration > Fabric > Nodes > User Management** to find the API details.

User	Scenario	UI	CLI	API
	Reset audit's expired password (Requires existing password)	Contact your administrator.	SSH or log in to the NSX CLI. Requires existing password. See Manage Local User's Password or Name Using the CLI	Go to the bios and update password.
Guest users	Reset your own forgotten password	Contact your administrator.	Contact your administrator.	Contact your administrator.
	Reset your own password (Requires known password)	If you know the guest user password, see Manage Local User Accounts .		Go to https://code.vmware.com/apis/1163/nsx-t and select your release number. Search for **System Administration > Configuration > Fabric > Nodes > User Management** to find the API details.
	Reset guest user's expired password	Contact Network Admin or administrator.	Contact your Network Admin or administrator.	Contact your Network Admin or administrator.

Note If you are an LDAP user, contact the LDAP administrator to reset your password.

Resetting Expired Passwords

In NSX, an admin user can reset your expired password. You must enter the expired user password to complete the password reset in the CLI.

Each user password has an expiration date tracked by NSX. If you log in with an expired password, the error message, `Your password has expired`, appears. A notification also appears if your password is expiring soon. Depending on what release you are using and what user you are, there are different actions to take.

The admin user can reset the expired passwords for their own, as well as the audit and Enterprise Admin user, using the CLI. The admin user can reset the expired passwords for all other local users, but not themselves in the UI.

When the admin resets any local user's expired password using the CLI and UI, the password expiration date does not reset. To extend the password expiration date to a new date or to the default 90 days, the admin must change it using the CLI or API. Another quick way for the admin to reset a local user's password expiration date is to deactivate and then activate the user. For details, see Manage Local User's Password or Name Using the CLI or Manage Local User Accounts .

Procedure

1 Choose one of the following steps depending on your release and needs:

- CLI - To reset the admin, Enterprise Admin, or audit user's expired passwords, run the `set user <username> password` command on the appliance's CLI as admin. You must know the current user password to complete the password reset. Guest users cannot be reset using the CLI.
- UI - The admin user can reset an expired password for other local users in the UI. Admin cannot change their own expired password using the UI. Look for the Notification message about the user password expiration and click on **Change Password**. Admin can also use **System > User Management** to change the passwords. Note, the password expiration date does not reset.
- API - Admin can go to https://developer.vmware.com/apis/1248/nsx-t and select your release number. Search for **System > Configuration > Fabric > Nodes > User Management** to find the API details.

2 Once the password reset is complete, the user gets prompted to change their password during login.

Resetting the Passwords of an Appliance

The following procedure applies to NSX Manager, NSX Edge, and Cloud Service Manager appliances.

Note If you have an NSX Manager cluster, resetting the password for any of the local users on one automatically resets the password for the other NSX Managers in the cluster. The synchronization of the password can take a few minutes.

If you have renamed the local user, use the new name in the following procedures.

When you reboot an appliance, the GRUB boot menu does not appear by default. Before you perform this procedure, you must configure GRUB so that the GRUB boot menu displays. For more information about configuring GRUB and changing the GRUB `root` password, see "Configure an Appliance to Display the GRUB Menu at Boot Time" in the *NSX Installation Guide*.

If you know the password for `root` but have forgotten the password for your local users, you can reset it using the following procedure:

1 Log in to the appliance as **root**.

2 To stop the server:

a For NSX Edge, run the command `/etc/init.d/nsx-edge-api-server stop`.

b For a Cloud Service Manager, skip this step.

c Otherwise, run the command `/etc/init.d/nsx-mp-api-server stop`.

3 (Optional) To reset the password for **admin**, run the command `passwd admin`.

4 (Optional) To reset the password for **audit**, run the command `passwd audit`.

5 (Optional) To reset a guest user password, run the command `passwd guestusername`.

6 Run the command `touch /var/vmware/nsx/reset_cluster_credentials`.

7 To restart the server:

   a For NSX Edge, run the command `/etc/init.d/nsx-edge-api-server start`.

   b Otherwise, run the command `/etc/init.d/nsx-mp-api-server start`.

If you have forgotten the root user's password, you can reset it using the following procedure. You can then use the above procedure to reset the password for all local users.

Procedure

1 Connect to the console of the appliance.

2 Reboot the system.

3 When the GRUB boot menu appears, press the left **SHIFT** or **ESC** key quickly. If you wait too long and the boot sequence does not pause, you must reboot the system again.

4 Press **e** to edit the menu.

   Choose the top Ubuntu line then enter the user name **root** and the GRUB password for root (not the same as the appliance's user root). The default password is `NSX@VM!WaR10`.

5 Press **e** to edit the selected option.

6 Search for the line starting with `linux` and add `systemd.wants=PasswordRecovery.service` to the end of the line.

7 Press **Ctrl-X** to boot.

8 When the log messages stop, enter the new password for root.

9 Enter the password again.

   The boot process continues.

10 After the reboot, you can verify the password change by logging in as root with the new password.

Authentication Policy Settings

You can view or customize your authentication policy settings through the API and CLI.

Some details are important to understand about viewing or changing policy settings.

- This feature allows one password policy only.
- Changes to the password configuration affects new users immediately. After administrators update existing user password configurations current users must follow updated password configuration changes.

- API configuration changes take about 20 seconds to take effect.
- For NSX Edge, changed passwords are not synchronized across the cluster. Both CLI and API changes appear in one node.
- Appliance transport node support includes BCG, autonomous Edge, or Unified Appliances (UA). There is no support for transport node password configuration changes for the local administrator or the auditor user in ESX.
- Privileged Access Management (PAM) supports setting minimum or maximum password length, but not both settings.
- Use of negative numbers for password entries sets the minimum range while positive numbers set the maximum range. To keep the existing default, leave the response empty.
- If modified pre-upgrade, the password policy configuration remains the same post upgrade for existing users. NSX Manager does not enforce the default password policy in this case.
- Authentication policies

```
[API] /api/v1/node/aaa/auth-policy
```

```
[API] /api/v1/cluster/<Node-UUID>/node/aaa/auth-policy
```

```
[API] /api/v1/transport-nodes/<transport-node-id>/node/aaa/auth-policy
```

NSX Manager includes the following CLI and API password complexity and authentication command support. These password policy options now sync across the management cluster nodes. Viewing password details does not require any permissions. Modifying existing password defaults requires admin permissions.

For more information on defaults ranges and other details, see the *NSX Command-Line Interface Reference* and the *NSX API Guide*.

Table 22-1. CLI Password Policy Customizable Options

Password Option	CLI Command

View or configure password complexity configuration	`get password-complexity`
	
	`Wed Jun 08 2022 UTC 12:57:45.325`
	`- minimum 12 characters in length`
	`- maximum 128 characters in length`
	`- minimum 1 lowercase characters`
	`- minimum 1 uppercase characters`
	`- minimum 1 numeric characters`
	`- minimum 1 special characters`
	`- default password complexity rules as enforced by the Linux PAM module`
	
	`set password-complexity`
	
	You can change specific parameters using these arguments:
	
	`Minimum password length (leave empty to not change):`
	`Maximum password length (leave empty to not change):`
	`Lower characters (leave empty to not change):`
	`Upper characters (leave empty to not change):`
	`Numeric characters (leave empty to not change):`
	`Special characters (leave empty to not change):`
	`Minimum unique characters (leave empty to not change):`
	`Allowed similar consecutives (leave empty to not change):`
	`Allowed monotonic sequence (leave empty to not change):`
	`Hash algorithm (leave empty to not change):`
	`Password remembrance (leave empty to not change):`

View authentication policy	`get auth-policy cli`
	
	`lockout-period   Lockout`

Table 22-1. CLI Password Policy Customizable Options (continued)

Password Option	CLI Command
	period max-auth-failures Maximum authentication failures before lockout
Configure authentication policy	set auth-policy cli lockout-period Lockout period max-auth-failures Maximum authentication failures before lockout

Table 22-2. API Password Policy Customizable Options

Password Option	API Commands

View password complexity and authorization policy	get /api/v1/node/aaa/auth-policy
	
	{
	 "_retry_prompt": 3,
	 "_schema": "AuthenticationPolicyProperties",
	 "_self": {
	 "href": "/node/aaa/auth-policy",
	 "rel": "self"
	 },
	 "api_failed_auth_lockout_period": 5,
	 "api_failed_auth_reset_period": 900,
	 "api_max_auth_failures": 900,
	 "cli_failed_auth_lockout_period": 900,
	 "cli_max_auth_failures": 5,
	 "digits": -1,
	 "hash_algorithm": "sha512",
	 "lower_chars": -1,
	 "max_repeats": 0,
	 "max_sequence": 0,
	 "maximum_password_length": 128,
	 "minimum_password_length": 12,
	 "minimum_unique_chars": 0,
	 "password_remembrance": 0,
	 "special_chars": -1,
	 "upper_chars": -1
	}

View VMware Identity Manager(VIDM) authorization policy	get auth-policy vidm Wed Jun 08 2022 UTC 12:58:28.357 LB Enabled: False Enabled: False Hostname: Thumbprint: Client Id: Node Hostname:

Table 22-2. API Password Policy Customizable Options (continued)

Password Option	API Commands

Configure VIDM authorization policy	`set auth-policy vidm`
	
	`enabled  Enabled property`
	` hostname  System's network name`
	` lb-extern External Load Balancer Flag For vIDM Wiring`

Configure password complexity and authorization policy
■ API failed authorization attempts lockout period (in seconds)
	`put /api/v1/node/aaa/auth-policy`
	
	`lockout_period <lockout-period-arg>`

■ Authentication failure lockout reset period	`lockout_reset_period`
■ Number of failures allowed before API lockout	`max_auth_failures`
■ Number of numeric characters	`digits`
■ Hash algorithm	`hash_algorithm`
■ Number of lowercase characters	`lower_chars`
■ Sequence of same characters	`max_repeats`
■ Maximum monotonic character sequence (1234 or DCBA)	`max_sequence`
■ Maximum password length	`maximum_password_length`
■ Minimum password length	`minimum_password_length`
■ Minimum unique characters	`minimum_unique_chars`
■ Minimum number of password reuse	`password_remembrance` ■ If 0, the check for previous passwords is disabled and users can reuse any previous password. Default. ■ If you enter a number, the user cannot reuse that number of previous passwords. For example, if set to 2, the user cannot reuse the last two passwords.
■ Number of special characters	`special_chars`

Table 22-2. API Password Policy Customizable Options (continued)

Password Option	API Commands
■ Number of uppercase characters	`upper_chars`

Reset password complexity, authentication policy, or both	For node, transport-nodes, and clusters:
	`reset-password-complexity`
	`reset-auth-policies`
	`reset-all`

Integration with VMware Identity Manager/Workspace ONE Access

You can configure NSX Manager to authenticate users using VMware Identity Manager (vIDM).

Note: The new product name for VMware Identity Manager is Workspace ONE Access.

Time Synchronization between NSX Manager, vIDM, and Related Components

For authentication to work correctly, NSX Manager, vIDM and other service providers such as Active Directory must all be time synchronized. This section describes how to time synchronize these components.

VMware Infrastructure

Follow the instructions in the following KB articles to synchronize ESXi hosts.

- https://kb.vmware.com/kb/1003736
- https://kb.vmware.com/kb/2012069

Third-Party Infrastructure

Follow the vendor's documentation on how synchronize VMs and hosts.

Configuring NTP on the vIDM Server (Not Recommended)

If you are not able to synchronize time across the hosts, you can disable synchronizing to host and configure NTP on the vIDM server. This method is not recommend because it requires the opening of UDP port 123 on the vIDM server

- Check the clock on the vIDM server and make sure it is correct.

```
# hwclock
Tue May  9 12:08:43 2017  -0.739213 seconds
```

- Edit `/etc/ntp.conf` and add the following entries if they don't exist.

```
server time.nist.gov
server pool.ntp.org
server time.is dynamic
restrict 192.168.100.0 netmask 255.255.255.0 nomodify notrap
```

- Open UDP port 123.

```
# iptables -A INPUT -p udp --dport 123 -j ACCEPT
```

  Run the following command to check that the port is open.

```
# iptables -L -n
```

- Start the NTP service.

```
/etc/init.d/ntp start
```

- Make NTP run automatically after a reboot.

```
# chkconfig --add ntp
# chkconfig ntp on
```

- Check that the NTP server can be reached.

```
# ntpq -p
```

  The `reach` column should not show 0. The `st` column should show some number other than 16..

Obtain the Certificate Thumbprint from a vIDM Host

Before you configure the integration of vIDM with NSX, you must get the certificate thumbprint from the vIDM host.

You must use OpenSSL version 1.x or higher for the thumbprint. On a vIDM host of version 3.3.2 or earlier, the command `openssl` might be running an older version of OpenSSL. In that case, you must use the command `openssl1`. This command is only available on a vIDM host.

You can check your version of OpenSSL with the following command:

```
openssl version
```

On a server that is not the vIDM host, you can use the `openssl` command that is running OpenSSL version 1.x or later.

Procedure

1 Log in at the vIDM host's console, or SSH to the vIDM host as the user **sshuser**, or log in to any server that can ping the vIDM host.

2 Run one of the following commands to get the thumbprint of the vIDM host.

- If you are logged in to a server that can ping the vIDM host, run the `openssl` command to get the thumbprint:

```
openssl s_client -connect <FQDN of vIDM host>:443 < /dev/null 2> /dev/null | openssl
x509 -sha256 -fingerprint -noout -in /dev/stdin
```

- If you are logged in to the vIDM host, do one of the following:
  - If the OpenSSL version is 0.9.x or earlier, run the following command:

```
openssl1 s_client -connect <FQDN of vIDM host>:443 < /dev/null 2> /dev/null |
openssl x509 -sha256 -fingerprint -noout -in /dev/stdin
```

    If you get an error running the command, you might need to run `openssl1` with the `sudo` command, that is, `sudo openssl1 ...`.

  - If the OpenSSL version is 1.x or later, run the following command:

```
openssl s_client -connect <FQDN of vIDM host>:443 < /dev/null 2> /dev/null |
openssl x509 -sha256 -fingerprint -noout -in /dev/stdin
```

    If you get an error running the command, you might need to run `openssl` with the `sudo` command, that is, `sudo openssl ...`.

Configure VMware Identity Manager/Workspace ONE Access Integration

You can integrate NSX with VMware Identity Manager (vIDM), which provides identity management services. The vIDM deployment can be a standalone vIDM host or a vIDM cluster.

Note: The new product name for VMware Identity Manager is VMware Workspace ONE Access.

The vIDM host or all the vIDM cluster components should have a certificate signed by a certificate authority (CA). Otherwise, logging in to vIDM from NSX Manager might not work with certain browsers, such as Microsoft Edge or Internet Explorer 11. For information about installing a CA-signed certificate on vIDM, see the VMware Identity Manager documentation at https://docs.vmware.com/en/VMware-Identity-Manager/index.html.

When you register NSX Manager with vIDM, you specify a redirect URI that points to NSX Manager. You can provide either the fully qualified domain name (FQDN) or the IP address. It is important to remember whether you use the FQDN or the IP address. When you try to log in to NSX Manager through vIDM, you must specify the host name in the URL the same way, that is, if you use the FQDN when registering the manager with vIDM, you must use the FQDN in the URL, and if you use the IP address when registering the manager with vIDM, you must use the IP address in the URL. Otherwise, login will fail.

If NSX API access is needed, one of the following configurations must be true:

- vIDM has a known CA-signed certificate.

- vIDM has the connector CA certificate trusted on the vIDM service side.
- vIDM uses outbound connector mode.

Note NSX Managers and vIDM must be in the same time zone. The recommended way is to use UTC.

You must configure your DNS servers to have PTR records if you are not using Virtual IP or an external load balancer (this means that the manager is configured using the physical IP or FQDN of the node).

If you configure vIDM to be integrated with an external load balancer, you must enable session persistence on the load balancer to avoid issues such as pages not loading or a user being unexpectedly logged out.

If the vIDM deployment is a vIDM cluster, the vIDM load balancer must be configured for SSL termination and re-encryption.

With vIDM enabled, you can still log in to NSX Manager with a local user account if you use the URL `https://<nsx-manager-ip-address>/login.jsp?local=true`.

If you use the UserPrincipalName (UPN) to log in to vIDM, authentication to NSX might fail. To avoid this issue, use a different type of credentials, for example, SAMAccountName.

If using NSX Cloud, you can log in to CSM separately using the URL `https://<csm-ip-address>/login.jsp?local=true`

Prerequisites

- Verify that you have the certificate thumbprint from the vIDM host or the vIDM load balancer, depending on the type of vIDM deployment (a standalone vIDM host or a vIDM cluster). The command to obtain the thumbprint is the same in both cases. See Obtain the Certificate Thumbprint from a vIDM Host.
- Verify that NSX Manager is registered as an OAuth client to vIDM. During the registration process, note the client ID and the client secret. For more information, see the VMware Identity Manager documentation at https://docs.vmware.com/en/VMware-Workspace-ONE-Access/3.3/idm-administrator/GUID-AD4B6F91-2D68-48F2-9212-5B69D40A1FAE.html. When you create the client, you only need to do the following:
  - Set **Access Type** to **Service Client Token**.
  - Specify a client ID.
  - Expand the **Advanced** field and click **Generate Shared Secret**.
  - Click **Add**.

  NSX Cloud Note If using NSX Cloud, also verify that CSM is registered as an OAuth client to vIDM.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **System > User Management > Authentication Providers > VMware Identity Manager**

3 Click **Edit**.

4 To enable external load balancer integration, click the **External Load Balancer Integration** toggle.

Note If you have Virtual IP (VIP) set up (check **System > Appliances > Virtual IP**), you cannot use the **External Load Balancer Integration** even if you enable it. This is because you can either have VIP or the External Load Balancer while configuring vIDM but not both. Disable VIP if you want to use the External Load Balancer. See Configure a Virtual IP (VIP) Address for a Cluster in the *NSX Installation Guide* for details.

5 To enable VMware Identity Manager integration, click the **VMware Identity Manager Integration** toggle.

6 Provide the following information.

Parameter	Description
VMware Identity Manager Appliance	The fully qualified domain name (FQDN) of the vIDM host or the vIDM load balancer, depending on the type of vIDM deployment (a standalone vIDM host or a vIDM cluster).
OAuth Client ID	The ID that is created when registering NSX Manager to vIDM.
OAuth Client Secret	The secret that is created when registering NSX Manager to vIDM.
SSL Thumbprint	The certificate thumbprint of the vIDM host. It must be an SHA-256 thumbprint.
NSX Appliance	The IP address or fully qualified domain name (FQDN) of NSX Manager. If you are using an NSX Manager cluster, use the load balancer FQDN or cluster VIP FQDN or IP address. If you specify a FQDN, you must access NSX Manager from a browser using the manager's FQDN in the URL, and if you specify an IP address, you must use the IP address in the URL. Alternatively, the vIDM administrator can configure the NSX Manager client so that you can connect using either the FQDN or the IP address.

7 Click **Save**.

8 If using NSX Cloud, repeat steps 1 through 8 from the CSM appliance by logging in to CSM instead of NSX Manager.

Validate VMware Identity Manager™ Functionality

After configuring VMware Identity Manager™, validate the functionality. Unless VMware Identity Manager™ is properly configured and validated, some users may receive Not Authorized (Error Code 98) messages when trying to log in.

Unless VMware Identity Manager™ is properly configured and validated, some users may receive Not Authorized (Error Code 98) messages when trying to log in.

Procedure

1 Create a base64 encoding of the username and password.

Run the following command to get the encoding and remove the trailing '\n' character. For example:

```
echo -n 'sfadmin@ad.node.com:password1234!' | base64 | tr -d '\n'
c2ZhZG1pbkBhZC5ub2RlLmNvbTpwYXNzd29yZDEyMzQhCg==
```

2 Verify that each user can make API call to each node.

Use a Remote Authorization curl command: `curl -k -H 'Authorization: Remote <base64 encoding string>' https://<node FQDN>/api/v1/node/aaa/auth-policy`. For example:

```
curl -k -H 'Authorization: Remote c2ZhZG1pbkBhZC5ub2RlLmNvbTpwYXNzd29yZDEyMzQhCg==' /
https://tmgr1.cptroot.com/api/v1/node/aaa/auth-policy
```

This returns the authorization policy settings, such as:

```
{
  "_schema": "AuthenticationPolicyProperties",
  "_self": {
    "href": "/node/aaa/auth-policy",
    "rel": "self"
  },
  "api_failed_auth_lockout_period": 900,
  "api_failed_auth_reset_period": 900,
  "api_max_auth_failures": 5,
  "cli_failed_auth_lockout_period": 900,
  "cli_max_auth_failures": 5,
  "minimum_password_length": 12
}
```

If the command does not return an error, the VMware Identity Manager™ is working correctly. No further steps are required. If the curl command returns an error, the user may be locked out.

Note Account lockout policies are set set and enforced on a per node basis. If one node in the cluster has locked out a user, other nodes may have not.

3 To reset a user lockout on a node:

a Retrieve the authorization policy using the local NSX Manager admin user:

```
curl -k -u 'admin:<password>' https://nsxmgr/api/v1/node/aaa/auth-policy
```

b Save the output to a JSON file in current working directory.

c Modify the file to change lockout period settings.

For example, many of the default settings apply lockout and reset periods of 900 seconds. Change these values to enable immediate reset, such as:

```
{
  "_schema": "AuthenticationPolicyProperties",
  "_self": {
    "href": "/node/aaa/auth-policy",
    "rel": "self"
  },
  "api_failed_auth_lockout_period": 1,
  "api_failed_auth_reset_period": 1,
  "api_max_auth_failures": 5,
  "cli_failed_auth_lockout_period": 1,
  "cli_max_auth_failures": 5,
  "minimum_password_length": 12
}
```

d Apply the change to the affected node.

```
curl -k -u 'admin:<password>' -H 'Content-Type: application/json' -d \
@<modified_policy_setting.json> https://nsxmgr/api/v1/node/aaa/auth-policy
```

e (Optional) Return the authorization policy settings files to its previous settings.

This should resolve the lockout issue. If you can still make remote auth API calls, but are still unable to log in through the browser, the browser may have an invalid cache or cookie stored. Clear your cache and cookies, and try again.

Integration with LDAP

You can configure NSX Manager to authenticate users using a directory service such as Active Directory over LDAP or OpenLDAP.

If you are using Active Directory (AD), and your AD forest is comprised of multiple subdomains, you should point NSX at your AD Global Catalog (GC) and configure each subdomain as an alternative domain name in NSX. The Global Catalog service usually runs on your primary AD domain controllers, and is a read-only copy of the most important information from all the primary and secondary domains. The GC service runs on port 3268 (plaintext), and 3269 (LDAP over TLS, encrypted).

For example, if your primary domain is "example.com" and you have subdomains "americas.example.com" and "emea.example.com", you should:

1 Configure NSX to use either the LDAP protocol on port 3268 or the LDAPS protocol on port 3269.

2 Add alternative domain names "americas.example.com" and "emea.example.com" in the NSX LDAP configuration.

Users in one of the subdomains must log in using the appropriate domain in their login name. For example, user "john" in the emea.example.com domain, must log in with the username "john@emea.example.com".

LDAP support on a Global Manager (NSX Federation) is identical to a Local Manager. LDAP configuration is not synchronized from Global Manager to Local Managers. Each NSX cluster should be configured separately for LDAP.

LDAP Identity Source

NSX Manager acts as an LDAP client, and interfaces with LDAP servers.

Three identity sources can be configured for user authentication. When a user logs into NSX Manager, the user is authenticated against the appropriate LDAP server of the user's domain. The LDAP server responds back with the authentication results, and the user group information. Once successfully authenticated, the user is assigned the roles corresponding to the groups that they belong to.

When integrating with Active Directory, NSX Manager allows users to log in using their samAccountName, or userPrincipalName. If the @domain portion of the userPrincipalName does not match the domain of the Active Directory instance, then you should also configure an alternative domain in the LDAP configuration for NSX.

In the following example, the domain of the Active Directory instance is "example.com" and a user with a samAccountName "jsmith" has a userPrincipalName of John.Smith@acquiredcompany.com. If you configure an alternative domain of "acquiredcompany.com", then this user can log in as "jsmith@example.com" using the samAccountName, or as John.Smith@acquiredcompany.com using the userPrincipalName. If the userPrincipalName has no @domain portion, the user won't be able to log in.

Logging in as jsmith@acquiredcompany.com will not work because the samAccountName can only used with the primary domain.

Procedure

1 Navigate to **System > User Management > LDAP**.

2 Click **Add Identity Source**.

3 Enter a **Name** for the identity source.

4 Enter the **Domain Name** This must correspond to the domain name of your Active Directory server, if using Active Directory.

5 Select the type: either **Active Directory over LDAP** or **Open LDAP**.

6 Click **Set** to configure LDAP servers. You can add up to three LDAP servers for failover support, to each domain.

Hostname/IP	The hostname or IP address of your LDAP server.
LDAP Protocol	Select the **protocol**: LDAP (unsecured) or LDAPS (secured).

Port	The default port populates based on the selected protocol. If your LDAP server runs on a non-standard port, you can edit this text box to give the port number.
Connection Status	Fill in the mandatory text boxes, including the LDAP server information, then click **Connection Status** to test the connection.
Use StartTLS	If selected, the LDAPv3 StartTLS extension is used to upgrade the connection to use encryption. To determine if you should use this option, consult your LDAP server administrator. This option is only available if LDAP protocol is selected.
Certificate	■ If you use LDAPS or LDAP + StartTLS, enter the PEM-encoded X.509 certificate of the server in the text box. If you leave this text box blank and click the **Check Status** link, NSX connects to the LDAP server. NSX then retrieves the LDAP server's certificate and asks if you want to trust that certificate. If you verify that the certificate is correct, click **OK**. The certificate text box gets populated with the retrieved certificate. ■ If your hostname/IP is an L4 Load Balancer VIP, the LDAP servers behind the VIP must present certificates signed by the same certificate authority (CA). You must enter the PEM-encoded X.509 certificate of the CA that signed the certificates. If you do not enter the certificate of the CA, NSX prompts you to accept the certificate of one of the LDAP servers, one which the load balancer randomly selects. If the server presents the full trust chain, including the certificate of the CA that signed the certificates of the other servers in the pool, the LDAP connection works when routed to another server. If the certificate initially presented does not include the CA certificate, the certificate presented by the other LDAP servers gets denied. For this reason, you must enter the certificate of the CA that signed all the certificates presented by the different LDAP or LDAPs servers. ■ If the LDAP servers are behind an L4 Load Balancer VIP, NSX will support certificates of the LDAP servers signed by different CAs if those CAs are subordinate to the same root CA. In this case, you must add the root CA certificate to the certificate field in the NSX LDAP configuration
Bind Identity	Enter the format as user@domainName, or you can specify the distinguished name. For Active Directory, use either the userPrincipalName (user@domainName) or the distinguished name. For OpenLDAP, you must supply a distinguished name. This text box is required unless your LDAP server supports anonymous bind, then it is optional. Consult your LDAP server administrator if you are not sure.
Password	Enter a password for the LDAP server. This text box is required unless your LDAP server supports anonymous bind, then it is optional. Consult your LDAP server administrator.

7 Click **Add**.

8 Enter the **Base DN**.

To add an Active Directory domain, a base distinguished name (Base DN) is needed. A Base DN is the starting point that an LDAP server uses when searching for users authentication within an Active Directory domain. For example, if your domain name is corp.local the DN for the Base DN for Active Directory is "DC=corp,DC=local".

All of the user and group entries you intend to use to control access to NSX must be contained within the LDAP directory tree rooted at the specified Base DN. If the Base DN is set to something too specific, such as an Organizational Unit deeper in your LDAP tree, NSX may not be able to find the entries it needs to locate users and determine group membership. Selecting a broad Base DN is a best practice if you are unsure.

9 Your NSX end users can now log in using their login name followed by @ and the domain name of your LDAP server, **user_name@domain_name**.

What to do next

Assign roles to users and groups. See Add a Role Assignment or Principal Identity.

Configuring Both vIDM and LDAP or Transitioning from vIDM to LDAP

If you have configured vIDM as the authentication server, you can add LDAP as an additional authentication server. You can also disable vIDM and use LDAP exclusively.

To configure vIDM integration, see Integration with VMware Identity Manager/Workspace ONE Access . To configure LDAP integration, see Integration with LDAP.

If you have both vIDM and LDAP integration configured, the URL for the login page for vIDM users is `https://<nsx-manager-ip-address>`. Users will be redirected to the vIDM login page. The URL for the login page for LDAP users is `https://<nsx-manager-ip-address>/login.jsp?local=true` and the login name must be in the format `user_name@domain_name`.

If you only have LDAP integration configured, the URL for the login page for vIDM users is `https://<nsx-manager-ip-address>` and the login name must be in the format `user_name@domain_name`.

If you have vIDM integration configured and want to transition to using LDAP only, first configure LDAP integration. The AD servers must be the same as the AD servers used in vIDM. Then disable vIDM on the vIDM configuration page. The roles, users, and role assignments created in vIDM will exist in LDAP.

NSX API Authentication Using a Session Cookie

This topic describes how to use NSX session-based authentication to generate a JSESSIONID cookie when using the API. Use this method to reduce the number of times you have to enter

your username and password. You can use this type of authentication with vIDM and LDAP authentication.

NSX uses several different mechanisms to authenticate NSX users. They include:

- HTTP authentication
- Session-based authentication
- Principal identity or certificate-based authentication
- Single sign on using vIDM and Cloud Service Platform (CSP)

The NSX uses a username and password to generate a session cookie during session creation. Once the session cookie has been created, subsequent API requests might optionally use this session cookie instead of the user name and password credentials. This means that the session state is specific to the particular NSX Manager that generated the session cookie. When clients make requests to the NSX Manager, it only allows clients to authenticate if the session cookie they present matches one of the cookies generated by the manager. When any user logs out of NSX Manager, the session cookie is immediately eliminated from the reverse-proxy of the NSX Manager and cannot be reused. Idle sessions time out automatically or you can delete them using the API.

Access using the API request generates audit log details. This logging is always enabled and cannot be disabled. Auditing of sessions is initiated at system startup. Audit log messages include the text audit="true" in the structured data part of the log message.

This example describes using cURL to create session-based authentication for API calls.

Procedure

1 To create a new session cookie that authenticates to the NSX Manager and retrieves XSRF from header enter:

```
# curl -vvv -k -c session.txt -X POST
-d 'j_username=admin@mywork.com&j_password=SecretPwsd3c4d' https://172.182.183.135/api/
session/create
        2>&1 > /dev/null | grep XSRF
< X-XSRF-TOKEN: 5a764b19-5ad2-4727-974d-510acbc171c8
```

In this example, the cURL command authenticates to the manager, places the session cookie in the sessions.txt file and writes all HTTP response headers to the headers.txt file. You will need to use one of the headers from headers.txt, the X-XSRF-TOKEN header, to provide in subsequent requests.

You can also use the standard unicode/URI encoding for the @ in the username.

An example of the session contents follows:

```
# cat session.txt
      # Netscape HTTP Cookie File
      # https://curl.haxx.se/docs/http-cookies.html
      # This file was generated by libcurl! Edit at your own
        risk.
      # HttpOnly_172.182.183.135 FALSE / TRUE 0 JSESSIONID CFG588DF6DGF493C0EAEFC62685C42E1
```

2 If you need to create two sessions, change the name of the session.txt files so that both sessions are valid.

```
# curl -vvv -k -c session1.txt -X POST -d 'j_username=
admin@mywork.com&j_password=SecretPwsd3c4d' https://172.182.183.135/api/session/create
2>&1 > /dev/null | grep XSRF
        < X-XSRF-TOKEN: cbce48f3-48fc-46c0-a8e7-f2f55ebf8e15
# curl -vvv -k -c session2.txt -X POST -d 'j_username=
admin@mywork.com&j_password=SecretPwsd3c4d' https://172.182.183.135/api/session/create
2>&1 > /dev/null | grep XSRF
        < X-XSRF-TOKEN: abf1fa2c-86d5-47e4-9a0a-242424dd1761
```

3 For subsequent calls, use the session cookie and XSRF header from the previous step.

```
# curl -k -b session.txt -H "X-XSRF-TOKEN: 5a764b19-5ad2-4727-974d-510acbc171c8"
      https://10.182.183.135/policy/api/v1/infra/segments
      {
        "results" : [ {
          "type" : "ROUTED",
          "subnets" : [ {
            "gateway_address" : "192.168.10.1/24",
            "network" : "192.168.10.0/24"
          } ],
          "connectivity_path" :
        "/infra/tier-1s/test_t1",
          "transport_zone_path" :
        "/infra/sites/default/enforcement-points/default/transport-zones/1b3a2f36-
bfd1-443e-a0f6-4de01abc963e",
          "advanced_config" : {
            "address_pool_paths" : [ ],
            "hybrid" : false,
            "multicast" : true,
            "inter_router" : false,
            "local_egress" : false,
            "urpf_mode" : "STRICT",
            "connectivity" : "ON"
          },
          "admin_state" : "UP",
          "replication_mode" : "MTEP",
          "resource_type" : "Segment",
          "id" : "seg1",
          "display_name" : "seg1",
          "path" : "/infra/segments/seg1",
          "relative_path" : "seg1",
          "parent_path" : "/infra",
```

```
      "unique_id" :
    "6573d2c9-f4f9-4b37-b410-71bded8857c3",
      "marked_for_delete" : false,
      "overridden" : false,
      "_create_user" : "admin",
      "_create_time" : 1633331197569,
      "_last_modified_user" : "admin",
      "_last_modified_time" : 1633331252660,
      "_system_owned" : false,
      "_protection" : "NOT_PROTECTED",
      "_revision" : 1
    } ],
    "result_count" : 1,
    "sort_by" : "display_name",
    "sort_ascending" : true
```


If you use the same session cookie with another node in the cluster, the command fails with the error message:


```
The credentials were incorrect or the account specified has been locked.","error_code":403.
```


When the session expires, NSX Manager responds with a 403 Forbidden HTTP response. You must then obtain a new session cookie and X-XSRF-TOKEN.

4 To configure the session expiry setting, use the connection_timeout API command. The default session expiry is 1800 seconds (30 minutes).

a Use `GET https://<nsx-mgr>/api/v1/cluster/api-service` to view a snapshot of the current configuration..

```
GET https://<nsx-mgr>/api/v1/cluster/api-service
{
    "session_timeout": 1800,
    "connection_timeout": 30,
    "protocol_versions": [
        {
            "name": "TLSv1.1",
            "enabled": true
        },
        {
            "name": "TLSv1.2",
            "enabled": true
        }
    ],
    "cipher_suites": [
        {
            "name": "TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA",
            "enabled": true
        },
        {
            "name": "TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256",
            "enabled": true
        },
        {
            "name": "TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256",
            "enabled": true
        },
        {
            "name": "TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA",
            "enabled": true
        },
        {
            "name": "TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384",
            "enabled": true
        },
        {
            "name": "TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384",
            "enabled": true
        },
        {
            "name": "TLS_RSA_WITH_AES_128_CBC_SHA",
            "enabled": true
        },
        {
            "name": "TLS_RSA_WITH_AES_128_CBC_SHA256",
            "enabled": true
        },
        {
```

```
            "name": "TLS_RSA_WITH_AES_128_GCM_SHA256",
            "enabled": true
        },
        {
            "name": "TLS_RSA_WITH_AES_256_CBC_SHA",
            "enabled": true
        },
        {
            "name": "TLS_RSA_WITH_AES_256_CBC_SHA256",
            "enabled": true
        },
        {
            "name": "TLS_RSA_WITH_AES_256_GCM_SHA384",
            "enabled": true
        }       ,
        {
            "name": "TLS_ECDSA_WITH_AES_256_GCM_SHA384",
            "enabled": true
        }
    ],
    "redirect_host": "",
    "client_api_rate_limit": 100,
    "global_api_concurrency_limit": 199,
    "client_api_concurrency_limit": 40,
    "basic_authentication_enabled": true,
    "cookie_based_authentication_enabled": true,
    "resource_type": "ApiServiceConfig",
    "id": "reverse_proxy_config",
    "display_name": "reverse_proxy_config",
    "_create_time": 1658339081246,
    "_create_user": "system",
    "_last_modified_time": 1658339081246,
    "_last_modified_user": "system",
    "_system_owned": false,
    "_protection": "NOT_PROTECTED",
    "_revision": 0
}
```


b Edit the JSON with the desired values and remove all the underscore fields (_xxxxx), except for `_revision`.


```
 GET https://<nsx-mgr>/api/v1/cluster/api-service
{
    "global_api_concurrency_limit": 199,
    "client_api_rate_limit": 100,
    "client_api_concurrency_limit": 40,
    "connection_timeout": 30,
    "redirect_host": "",
    "cipher_suites": [
      {"enabled": true, "name": "TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384"},
      {"enabled": true, "name": "TLS_RSA_WITH_AES_256_GCM_SHA384"},
      {"enabled": true, "name": "TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256"},
      {"enabled": true, "name": "TLS_RSA_WITH_AES_128_GCM_SHA256"}
      {"enabled": true, "name": "TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384}",
```

```
      {"enabled": true, "name": "TLS_RSA_WITH_AES_256_CBC_SHA256"},
      {"enabled": true, "name": "TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA"},
      {"enabled": true, "name": "TLS_RSA_WITH_AES_256_CBC_SHA"},
      {"enabled": true, "name": "TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256"},
      {"enabled": true, "name": "TLS_RSA_WITH_AES_128_CBC_SHA256"},
      {"enabled": false, "name": "TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA"},
      {"enabled": true, "name": "TLS_RSA_WITH_AES_128_CBC_SHA"},

      {"enabled": true, "name": "TLS_ECDSA_WITH_AES_256_GCM_SHA384"},
    ],
      "protocol_versions": [
        {"enabled": true, "name": "TLSv1.1"},
    {"enabled": true, "name": "TLSv1.2"}
  ]
    "_revision": 0
}
```

c Use `PUT https://<nsx-mgr>/api/v1/cluster/api-service` to modify the current configuration.

5 To delete a session cookie, use the /api/session/destroy API command.

```
curl -k -b cookies.txt -H "`grep X-XSRF-TOKEN headers.txt`" -X POST https://<nsx-
manager>/api/session/destroy
```

For example:

```
curl -k -b cookies.txt -H "`grep X-XSRF-TOKEN headers.txt`" https://$TESTHOST/api/session/
destroy
curl -k -b cookies.txt -H "`grep X-XSRF-TOKEN headers.txt`" https://$TESTHOST/api/v1/
logical-ports
```

Response:

```
{
    "module_name" : "common-services",
    "error_message" : "The credentials were incorrect or the account specified has been
locked.",
    "error_code" : "403"
}
```

What to do next

To review the requirements to authenticate users with your session-based supported authentication service, see Integration with VMware Identity Manager/Workspace ONE Access or Integration with LDAP.

Add a Role Assignment or Principal Identity

You can assign roles to users or user groups if VMware Identity Manager™ is integrated with NSX, or if you have LDAP as an authentication provider. You can also assign roles to principal identities.

A principal is an component or a third-party application such as an OpenStack product. With a principal identity, a principal can use the identity name to create an object and ensure that only an entity with the same identity name can modify or delete the object. A principal identity has the following properties:

- Name
- Node ID - this can be any alphanumeric value assigned to a principal identity
- Certificate
- RBAC role indicating the access rights of this principal

Users (local, remote, or principal identity) with the Enterprise Administrator role can modify or delete objects owned by principal identities. Users (local, remote, or principal identity) without the Enterprise Administrator role cannot modify or delete protected objects owned by principal identities, but can modify or delete unprotected objects.

If a principal identity user's certificate expires, you must import a new certificate and make an API call to update the principal identity user's certificate (see the procedure below). For more information about the NSX API, a link to the API resource is available at https://docs.vmware.com/en/VMware-NSX/index.html.

A principal identity user's certificate must satisfy the following requirements:

- SHA256 based.
- RSA/DSA message algorithm with 2048 bits or above key size.
- Cannot be a root certificate.

You can delete a principal identity using the API. However, deleting a principal identity does not automatically delete the corresponding certificate. You must delete the certificate manually.

Steps to delete a principal identity and its certificate:

1 Get the details of the principal identity to delete and note the `certificate_id` value in the response.

```
GET /api/v1/trust-management/principal-identities/<principal-identity-id>
```

2 Delete the principal identity.

```
DELETE /api/v1/trust-management/principal-identities/<principal-identity-id>
```

3 Delete the certificate using the `certificate_id` value obtained in step 1.

```
DELETE /api/v1/trust-management/certificates/<certificate_id>
```

For LDAP, you configure user groups to user roles mapping information; the groups correspond to the user groups specified in the Active Directory (AD). To grant a user permissions on NSX, add that user to the mapped group in AD.

Prerequisites

You must have an authentication provider configured:

- For role assignment for vIDM, verify that a vIDM host is associated with NSX. For more information, see Configure VMware Identity Manager/Workspace ONE Access Integration.
- For role assignment for LDAP, verify that you have an LDAP identity source. For more information, see LDAP Identity Source.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Select **System > User Management**.
3. To assign roles to users, select **Add > Role Assignment for vIDM**.
   a. Select a user or user group.
   b. Select a role.
   c. Click **Save**.
4. To add a principal identity, select **Add > Principal Identity with Role**.
   a. Enter a name for the principal identity.
   b. Select a role.
   c. Enter a node ID.
   d. Enter a certificate in PEM format.
   e. Click **Save**.
5. To add a role assignment for LDAP select **Add > Role Assignment for LDAP**.
   a. Select a domain.
   b. Enter the first few characters of the user's name, login ID, or a group name to search the LDAP directory, then select a user or group from the list that appears.
   c. Select a role.
   d. Click **Save**.
6. (Optional) If using NSX Cloud, log in to the CSM appliance instead of NSX Manager and repeat steps 1 through 4.

7 If the certificate for the principal identity expires, perform the following steps:

a Import a new certificate and note the certificate's ID. See Import a Self-signed or CA-signed Certificate.

b Call the following API to get the ID of the principal identity.

```
GET https://<nsx-mgr>/api/v1/trust-management/principal-identities
```

c Call the following API to update the principal identity's certificate. You must provide the imported certificate's ID and the principal identity user's ID.

For example,

```
POST https://<nsx-mgr>/api/v1/trust-management/principal-identities?
action=update_certificate
{
    "principal_identity_id": "ebd3032d-728e-44d4-9914-d4f81c9972cb",
    "certificate_id" : "abd3032d-728e-44d4-9914-d4f81c9972cc"
}
```

Role-Based Access Control

With role-based access control (RBAC), you can restrict system access to authorized users. Users are assigned roles and each role has specific permissions.

There are four types of permissions.

- Full access (FA includes Create, Read, Update, and Delete)
- Execute (Read, Update)
- Read
- None

Full access gives the user all permissions.

NSX has the following built-in roles. Role names in the UI can be different in the API. In NSX, if you have permission, you can clone an existing role, add a new role, edit newly created roles, or delete newly created roles.

- Auditor
- Cloud Admin (Available in the Cloud environment only.)
- Cloud Operator (Available in the Cloud environment only.)
- Enterprise Admin
- GI Partner Admin (Guest Introspection)
- LB Admin (Load Balancer)
- LB Operator
- Network Admin

- Network Operator
- Netx Partner Admin (Network Introspection)
- Project Admin
- Security Admin
- Security Operator
- Support Bundle Collector
- VPN Admin

To view the built-in and custom roles and their associated permissions, navigate to **System > User Management > Roles** and expand the row to view details. You can view permissions of all categories by expanding the role's Permissions details.

After you have assigned an Active Directory (AD) user a role, if the username is changed on the AD server, you need to assign the role again using the new username.

The Roles and Permissions tables do not include the Project Admin role. The Project Admin role has full access to all configurations in a project. The Network Admin, Network Operator, Security Admin, and Security Operator roles in a project have RBAC permissions only within the scope of the project, and not for the entire NSX system.

Note For VMware NSX® Intelligence™ RBAC information, see the *Using and Managing VMware NSX Intelligence* documentation.

Roles and Permissions

The following tables, Table 22-3. Roles and Permissions and Table 22-4. Roles and Permissions for Manager Mode, show the permissions each role has for different operations. The following abbreviations are used:

- A - Auditor
- CA - Cloud Admin (Available in the Cloud environment only.)
- CO - Cloud Operator (Available in the Cloud environment only.)
- EA - Enterprise Admin
- GIPA - GI Partner Admin (Guest Introspection Partner Administrator)
- LBA - LB Admin (Load Balancer)
- LBO - LB Operator
- NA - Network Admin
- NO - Network Operator
- NXPA - Netx Partner Admin (Network Introspection Partner Administrator)
- SA - Security Admin

- SO - Security Operator
- SBC - Support Bundle Collector
- VPNA - VPN Admin
- FA - Full access
- E - Execute
- R - Read

Table 22-3. Roles and Permissions

Operation	EA	A	NA	NO	SA	SO	CA	CO	LBA	LBO	VPNA	GIPA	NXPA	SBC
Networking > Tier-0 Gateways	FA	R	FA	R	R	R	FA	R	R	R	R	R	R	None
Networking > Tier-1 Gateways	FA	R	FA	R	R	R	FA	R	R	R	R	R	R	None
Networking > Network Interface	FA	R	FA	FA	R	R	FA	R	R	R	R	R	R	None
Networking > Network Static Routes	FA	R	FA	FA	R	R	FA	R	R	R	R	R	R	None
Networking > Locale Services	FA	R	FA	FA	R	R	FA	R	R	R	R	R	R	None
Networking > Static ARP Configuration	FA	R	FA	FA	R	R	FA	R	R	R	R	R	R	None
Networking > Segments	FA	R	FA	R	R	R	FA	R	R	R	R	R	R	None

Table 22-3. Roles and Permissions (continued)

Operation	EA	A	NA	NO	SA	SO	CA	CO	LBA	LBO	VPNA	GIPA	NXPA	SBC
Networking > Segments > Segment Profiles	FA	R	FA	R	R	R	FA	R	R	R	R	R	R	None
Networking > IP Address Pools	FA	R	FA	FA	R	R	FA	R	R	R	None	None	None	None
Networking Forwarding Policies	FA	R	FA	R	FA	R	FA	R	None	None	None	None	None	None
Networking > DNS	FA	R	FA	FA	R	R	FA	R	R	R	None	None	None	None
Networking > DHCP	FA	R	FA	R	R	R	FA	R	R	R	None	None	None	None
Networking > Load Balancing	FA	R	None	None	R	None	FA	R	FA	R	None	None	None	None
Networking > NAT	FA	R	FA	R	FA	R	FA	R	R	R	None	None	None	None
Networking > VPN	FA	R	FA	R	FA	R	FA	R	None	None	FA	None	None	None
Networking > IPv6 Profiles	FA	R	FA	R	R	R	FA	R	R	R	None	None	None	None
Security > Distributed Firewall	FA	R	R	R	FA	R	FA	R	R	R	R	R	R	None

Table 22-3. Roles and Permissions (continued)

Operation	EA	A	NA	NO	SA	SO	CA	CO	LBA	LBO	VPNA	GIPA	NXPA	SBC
Security > Gateway Firewall	FA	R	R	R	FA	R	FA	R	None	None	None	None	FA	None
Security > Network Introspection	FA	R	R	R	FA	R	FA	R	None	None	None	None	FA	None
Security > Endpoint Protection Rules	FA	R	R	R	FA	R	FA	R	None	None	None	FA	None	None
Inventory > Context Profiles	FA	R	FA	R	FA	R	FA	R	R	R	R	R	R	None
Inventory > Virtual Machines	R	R	R	R	R	R	R	R	R	R	R	R	R	None
Inventory > Virtual Machines > Create & Assign Tags to VM	FA	R	R	R	FA	R	FA	R	R	R	R	FA	FA	None
Inventory > Containers	FA	R	R	R	R	R	None	None	None	None	None	None	None	None
Inventory > Physical Servers	FA	R	R	R	R	R	R	R	R	R	None	None	None	None

Table 22-3. Roles and Permissions (continued)

Operation	EA	A	NA	NO	SA	SO	CA	CO	LBA	LBO	VPNA	GIPA	NXPA	SBC
Plan & Troubleshoot > Port Mirroring	FA	R	FA	R	R	R	FA	R	None	None	None	None	None	None
Plan & Troubleshoot > Port Mirroring Binding	FA	R	FA	FA	R	R	FA	R	R	R	R	R	R	None
Plan & Troubleshoot > Monitoring Profile Binding	FA	R	FA	FA	R	R	FA	R	R	R	R	R	R	None
Plan & Troubleshoot > IPFIX > Firewall IPFIX Profiles	FA	R	FA	R	FA	R	FA	R	R	R	R	R	R	None
Plan & Troubleshoot > IPFIX > Switch IPFIX Profiles	FA	R	FA	R	R	R	FA	R	R	R	R	R	R	None
Plan & Troubleshoot > Traceflow	FA	FA	FA	FA	FA	FA	FA	FA	FA	FA	None	None	None	None
System > Fabric > Hosts > Clusters	FA	R	R	R	R	R	R	R	R	R	None	None	None	R

Table 22-3. Roles and Permissions (continued)

Operation	EA	A	NA	NO	SA	SO	CA	CO	LBA	LBO	VPNA	GIPA	NXPA	SBC
System > Fabric > Hosts > Other Nodes	FA	R	R	R	R	R	R	R	R	R	None	None	None	R
System > Fabric > Hosts > Standalone	FA	R	R	R	R	R	R	R	R	R	None	None	None	R
System > Fabric > Hosts > Transport Node Profile	FA	R	R	R	R	R	R	R	R	R	None	None	None	R
System > Fabric >Nodes Hosts	FA	R	R	R	R	R	R	R	None	None	None	None	None	None
System > Fabric > Nodes	FA	R	FA	R	FA	R	R	R	R	R	None	None	None	None
System > Fabric > Nodes > Edge Transport Nodes	FA	R	FA	R	R	R	R	R	None	None	None	None	None	None
System > Fabric > Nodes > Edge Clusters	FA	R	FA	R	R	R	R	R	None	None	None	None	None	None
System > Fabric > Nodes > Container Clusters	FA	R	FA	R	R	R	None	None	R	R	None	None	None	None

Table 22-3. Roles and Permissions (continued)

Operation	EA	A	NA	NO	SA	SO	CA	CO	LBA	LBO	VPNA	GIPA	NXPA	SBC
System > Fabric > Nodes > Transport Nodes	FA	R	R	R	R	R	R	R	R	R	None	None	None	R
System > Fabric > Nodes > Tunnels	R	R	R	R	R	R	R	R	R	R	None	None	None	None
System > Fabric > Profiles > Uplink Profiles	FA	R	R	R	R	R	R	R	R	R	None	None	None	None
System > Fabric > Profiles > Edge Cluster Profiles	FA	R	FA	R	R	R	R	R	R	R	None	None	None	None
System > Fabric > Profiles > Configuration	FA	R	None	None	None	None	R	R	None	None	None	None	None	None
System > Fabric > Profiles > Node Profiles	FA	R	R	R	R	R	R	R	R	R	None	None	None	None
System > Fabric > Transport Zones > Add Zones	FA	R	R	R	R	R	R	R	R	R	None	None	None	None

Table 22-3. Roles and Permissions (continued)

Operation	EA	A	NA	NO	SA	SO	CA	CO	LBA	LBO	VPNA	GIPA	NXPA	SBC
System > Fabric > Transport Zones > Transport Zone Profiles	FA	R	R	R	R	R	R	R	None	None	None	None	None	None
System > Fabric > Transport Zones > Health Configuration	FA	R	R	R	R	R	R	R	R	R	None	None	None	None
System > Fabric > Compute Managers	FA	R	R	R	R	R	R	R	None	None	None	R	R	None
System > Fabric > Settings	FA	None	None	None	None	None	None	None	None	None	None	None	None	None
System > Certificates	FA	R	None	None	FA	R	None	None	FA	R	FA	None	None	None
System > Service Deployments > Service Instances	FA	R	R	R	FA	R	FA	R	None	None	None	FA	FA	None
System > Support Bundle	FA	None	None	None	None	None	None	None	None	None	None	None	None	FA
System > Backup	FA	R	None	None	None	None	None	None	None	None	None	None	None	None

Table 22-3. Roles and Permissions (continued)

Operation	EA	A	NA	NO	SA	SO	CA	CO	LBA	LBO	VPNA	GIPA	NXPA	SBC
System > Restore	FA	R	None	None	None	None	None	None	None	None	None	None	None	None
System > Upgrade	FA	R	R	R	R	R	None	None	None	None	None	None	None	None
System > Migrate	FA	None	None	None	None	None	None	None	None	None	None	None	None	FA
System > User Mgt > User Role Assignments	FA	R	None	None	None	None	None	None	None	None	None	None	None	None
System > Local Users	FA	R	None	None	None	None	None	None	None	None	None	None	None	None
System > Roles	FA	R	FA	R	FA	FA	R	R	R	R	R	R	R	None
System > Authentication Providers	FA	R	FA	R	FA	FA	R	R	R	R	R	R	R	None
System > Licenses	FA	R	R	R	R	R	None	None	None	None	None	None	None	None
System > System Administration	FA	R	R	R	R	R	R	R	None	None	None	None	None	None
Custom Dashboard Configuration	FA	R	R	R	R	R	FA	R	R	R	R	R	R	None

Table 22-4. Roles and Permissions for Manager Mode

Operation	EA	A	NA	NO	SA	SO	CA	CO	LBA	LBO	VPNA	GIPA	NXPA	SBC
Plan & Troubleshoot > Port Connection	E	R	E	E	E	E	E	R	E	E	None	None	None	None
Plan & Troubleshoot > Traceflow	FA	R	E	E	E	E	None	None	E	E	None	None	None	None
Plan & Troubleshoot > Port Mirroring	FA	R	FA	R	R	R	FA	R	None	None	None	None	None	None
Plan & Troubleshoot > IPFIX	FA	R	FA	R	FA	R	FA	R	R	R	R	R	R	None
Security > Distributed Firewall > General	FA	R	R	R	FA	R	FA	R	None	None	None	None	R	None
Security > Distributed Firewall > Configuration	FA	R	R	R	FA	R	FA	R	None	None	None	None	None	None
Security > Edge Firewall	FA	R	R	R	FA	R	FA	R	None	None	None	None	FA	None
Networking > Routers	FA	R	FA	FA	R	R	FA	R	R	R	R	None	R	None
Networking > NAT	FA	R	FA	R	FA	R	FA	R	R	R	None	None	None	None

Table 22-4. Roles and Permissions for Manager Mode (continued)

Operation	EA	A	NA	NO	SA	SO	CA	CO	LBA	LBO	VPNA	GIPA	NXPA	SBC
Networking > DHCP > Server Profiles	FA	R	FA	R	None	None	FA	R	None	None	None	None	None	None
Networking > DHCP > Servers	FA	R	FA	R	None	None	FA	R	None	None	None	None	None	None
Networking > DHCP > Relay Profiles	FA	R	FA	R	None	None	FA	R	None	None	None	None	None	None
Networking > DHCP > Relay Services	FA	R	FA	R	None	None	FA	R	None	None	None	None	None	None
Networking > DHCP > Metadata Proxies	FA	R	FA	R	None	None	None	None	None	None	None	None	None	None
Networking > IPAM	FA	R	FA	FA	R	R	None	None	R	R	None	None	None	None
Networking > Logical Switches > Switches	FA	R	FA	R	R	R	FA	R	R	R	R	None	R	None
Networking > Logical Switches > Ports	FA	R	FA	R	R	R	FA	R	R	R	R	None	R	None

Table 22-4. Roles and Permissions for Manager Mode (continued)

Operation	EA	A	NA	NO	SA	SO	CA	CO	LBA	LBO	VPNA	GIPA	NXPA	SBC
Networking > Logical Switches > Switching Profiles	FA	R	FA	R	R	R	FA	R	R	R	None	None	None	None
Networking > Load Balancing > Load Balancers	FA	R	None	None	R	None	FA	R	FA	R	None	None	None	None
Networking > Load Balancing > Profiles > SSL Profiles	FA	R	None	None	FA	R	FA	R	FA	R	None	None	None	None
Inventory > Groups	FA	R	FA	R	FA	R	FA	R	R	R	R	R	R	None
Inventory > Groups > IP Sets	FA	R	FA	R	FA	R	FA	R	R	R	R	R	R	None
Inventory > IP Pools	FA	R	FA	R	None	None	None	None	R	R	R	R	R	None
Inventory > Groups > MAC Sets	FA	R	FA	R	FA	R	FA	R	R	R	R	R	R	None
Inventory > Services	FA	R	FA	R	FA	R	FA	R	R	R	R	R	R	None

Table 22-4. Roles and Permissions for Manager Mode (continued)

Operation	EA	A	NA	NO	SA	SO	CA	CO	LBA	LBO	VPNA	GIPA	NXPA	SBC
Inventory > Virtual Machines	R	R	R	R	R	R	R	R	R	R	R	R	R	None
Inventory > Virtual Machines > Create & Assign Tags to VM	FA	R	R	R	FA	R	FA	R	R	R	R	FA	FA	None
Inventory > Virtual Machines > Configure Tags	FA	None	None	None	None	None	None	None	None	None	None	None	None	None
System > Support Bundle	FA	None	None	None	None	None	None	None	None	None	None	None	None	FA

Create or Manage Custom Roles

Extend the RBAC capabilities provided by NSX and create custom roles that suit your operational requirements. You can clone an existing role and customize it or you can create a role afresh. You can also edit and delete user-created roles.

- You can create custom roles only for features available in the Policy mode. If you clone a role with access to features in the Manager mode, the cloned role provides access only to the Policy mode features. For example, features like Upgrade, Migrate, Fabric, TraceFlow, NSX Intelligence, and Inventory of Physical Servers and Containers are only available in Manager mode and therefore not supported. Most features are supported. The unsupported features for users with a custom role include:
  - **System > Configuration > Fabric > Profiles**
  - **System > Configuration > Fabric > Transport Zones**
  - **System > Configuration > Fabric > Settings > Tunnel/Remote and Tunnel Endpoint**

  - **System > Configuration > Identity Firewall AD**
  - **System > Lifecycle Management > Upgrade and Migrate**
  - **System > Settings > User Management, Support Bundle, Proxy Settings, and User Interface Settings**

  For more information on the Manager and Policy modes, see Chapter 1 NSX Manager.

- Only an Enterprise Administrator can assign the role management feature's permission to a custom role. An Enterprise Administrator can create a custom role to delegate further custom role creation and user role assignment.
- A user assigned with a custom role can only create other custom roles with equal or lower permission sets. A user with a custom role cannot create or assign roles with permissions higher than their own.
- A user assigned with a custom role cannot modify or delete the role assigned to them.

Note Custom roles are not supported on Global Manager (Federation).

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **System > Users and Roles > Roles**.

3 Clone an existing role or create one.

  - To clone a role, click  for that role and select **Clone**. Enter a name for the cloned role and specify permissions as per your operational requirements.
  - To create a role, click **Add Role**. Enter a name for the role and update the permissions as per your operational requirements.

  Note Based on the features you select, NSX might suggest additional permissions for the new role definition to be valid. Review the recommendations and click **Apply**.

  When creating a custom role, NSX checks for feature interdependencies. The interdependency check ensures that the user has a minimum of read access to the additional features that are required for the role to be valid.

  For example, if a user creates a role with full access permissions to Gateway Firewall and the **None** access permission to the Networking Gateway feature, the role is invalid. NSX then suggests that the user assign at least read access to the additionally required Networking Gateway feature.

4 (Optional) Edit or delete a user-created role.

  - To edit a user-created role, for example, if you wanted to extend access, click  for that role and select **Edit**. Change the role name, description, and permissions as per your operational requirements.

- To delete a user-created role, for example, if it was for temporary access, click  for that role and select **Delete**.

Logging User Account Changes

Changes to a user's role assignment are automatically written to syslog and the audit log.

For more information about syslog and the audit log, see Log Messages and Error Codes.

An example of a log message when assigning a role to a vIDM user:

```
2020-09-24T16:05:51.244Z nsxmanager-14663974-1-CertKB-FS NSX 5519 - [nsx@6876 audit="true"
comp="nsx-manager" entId="e3c2af75-9d0f-4020-90cc-f2f00d6af255" level="INFO"
reqId="b27711c6-0590-4b39-b8b6-f0980a0597f0" subcomp="policy" update="true" username="admin"]
UserName="admin", ModuleName="AAA", Operation="CreateRoleBinding", Operation
status="success", New
value=[{"name":"test_AU@idfw.local","type":"remote_user","identity_source_type":"VIDM","roles"
:[{"role":"auditor"}],"id":"bba634c9-cfbd-4806-a831-e63ec195e1f9","_protection":"UNKNOWN"}]
```

An example of a log message when updating the role of a vIDM user:

```
2020-09-24T16:12:51.217Z nsxmanager-14663974-1-CertKB-FS NSX 5519 -
[nsx@6876 audit="true" comp="nsx-manager" entId="e3c2af75-9d0f-4020-90cc-f2f00d6af255"
level="INFO" reqId="973faed4-f4b5-443d-bd79-7d995c027183" subcomp="policy" update="true"
username="admin"] UserName="admin", ModuleName="AAA", Operation="UpdateRoleBinding",
Operation status="success", New value=["e3c2af75-9d0f-4020-90cc-f2f00d6af255"
{"name":"test_AU@idfw.local","type":"remote_user","identity_source_type":"VIDM","roles":
[{"role":"security_admin"}],"_protection":"UNKNOWN"}]
```

An example of a log message when assigning a role to an LDAP user:

```
2020-09-24T16:06:28.663Z nsxmanager-14663974-1-CertKB-FS NSX 5519 - [nsx@6876 audit="true"
comp="nsx-manager" entId="35e45569-6da6-4dcd-b4a1-75747cdd6cf8" level="INFO"
reqId="db27f4ae-25a7-4482-b3f4-49228d12960b" subcomp="policy" update="true" username="admin"]
UserName="admin", ModuleName="AAA", Operation="CreateRoleBinding", Operation
status="success", New
value=[{"name":"skrasner@airius.com","type":"remote_user","identity_source_type":"LDAP","ident
ity_source_id":"ldap","roles":[{"role":"auditor"}],"id":"dd8d3675-
c574-454b-975e-300b65462827","_protection":"UNKNOWN"}]
```

An example of a log message when updating the role of an LDAP user:

```
2020-09-24T16:12:37.449Z nsxmanager-14663974-1-CertKB-FS NSX 5519 - [nsx@6876 audit="true"
comp="nsx-manager" entId="35e45569-6da6-4dcd-b4a1-75747cdd6cf8" level="INFO"
reqId="d7cdd3de-75a1-4d29-9fea-27e1dda4b5e2" subcomp="policy" update="true" username="admin"]
UserName="admin", ModuleName="AAA", Operation="UpdateRoleBinding", Operation
status="success", New value=["35e45569-6da6-4dcd-b4a1-75747cdd6cf8"
{"name":"skrasner@airius.com","type":"remote_user","identity_source_type":"LDAP","identity_sou
rce_id":"ldap","roles":[{"role":"network_admin"}],"_protection":"UNKNOWN"}]
```

Certificates

23

After you install NSX, the manager nodes and cluster have self-signed certificates.

If you are using NSX Federation, additional certificates are set up to establish trust between the Local Managers and Global Manager. If you are using TLS Inspection, a certificate authority (CA) security certificate is required. For details on TLS Inspection and certificates, see TLS Inspection.

You can import certificates, create a certificate signing request (CSR), generate self-signed certificates, and import a certificate revocation list (CRL). To improve security, it is recommended that you replace the self-signed certificates with CA-signed certificates.

This chapter includes the following topics:

- Types of Certificates
- Certificates for NSX and NSX Federation
- Create a Certificate Signing Request File
- Creating Self-signed Certificates
- Importing and Replacing Certificates
- Importing and Retrieving CRLs
- Import or Update a Trusted CA Bundle
- Storage of Public Certificates and Private Keys for Load Balancer or VPN service
- Alarm Notification for Certificate Expiration

Types of Certificates

There are three categories of self-signed certificates in NSX.

- Platform Certificates
- NSX Services Certificates
- Principal Identity Certificates

Refer to the following sections for details on each certificate category. Other types of certificates might be present in NSX. Refer to that component or application documentation for details on their certificates.

Starting in NSX and NSX Federation 4.1, you can replace the self-signed internal certificates with CA-signed certificates using ECDSA-based certificates that get generated using AES 256 encryption. For a list of replaceable certificates, refer to Certificates for NSX and NSX Federation. For command details, see the *NSX API Guide*.

Note Though NSX supports secp256k1 keys for all certificates, do not use this key if your environment requires only FIPS-approved cryptographic keys.

Platform Certificates

After you install NSX, navigate to **System > Certificates** to view the platform certificates created by the system. By default, these self-signed X.509 RSA 2048/SHA256 certificates get used for internal communication within NSX and for external authentication when you use NSX Manager to access the APIs or the UI.

The internal certificates are not viewable or editable.

If VMware Cloud Foundation™ (VCF) was used to deploy NSX, the default NSX API and Cluster certificates get replaced with CA certificates signed by the VMware Certificate Authority (VMCA) from vCenter. The API and Cluster certificates might still display in the certificate list, but are not used. Replace the CA-signed certificates using the procedure in the VCF Administration Guide. After you perform the replacement, your NSX Manager stores in the UI contain the API and Cluster certificates, the VMCA CA certificates, and the signed certificates by the third-party organization. From then on, the NSX Manager uses the signed certificate from your organization.

Table 23-1. Platform Certificates in NSX

Naming Convention in NSX Manager	Purpose	Replaceable?	Default Validity
`API (previously known as tomcat)`	This serves as a server certificate for the API/UI client reaching the NSX Manager HTTPS port (port 443).	Yes. See Replace Certificates	825 days
`Cluster (previously known as mp-cluster)`	This certificate serves as a server certificate for API/UI client reaching the HTTPS port of the cluster VIP (port 443) when a VIP is used for an NSX Manager cluster.	Yes. See Replace Certificates	825 days
Other certificates	Certificates specifically for other purposes including NSX Federation, Cluster Boot Manager (CBM), and Central Control Plane (CCP).	See Certificates for NSX and NSX Federation for details on self-signed certificates auto-configured for NSXand NSX Federation environments.	Varies

NSX Service Certificates

NSX service certificates are user-facing for services such as load balancer, VPN, and TLS Inspection. The policy API manages service certificates. Non-service certificates are used by the platform for tasks such as cluster management. The management plane (MP) or truststore APIs managed non-service certificates.

When adding service certificates using the policy API, the certificate gets sent to the MP/truststore API, but not with the relationship reversed.

NSX service certificates cannot be self signed. You must import them. See Importing and Replacing Certificates for instructions.

You can generate a root certificate authority (CA) certificate and a private key based on RSA. CA certificates are able to sign other certificates.

You can use a certificate signing request (CSR) as an NSX service certificate if it is signed by a CA (a local CA or a public CA like Verisign). Once the CSR is signed, you can import that signed certificate into NSX Manager. A CSR can be generated on NSX Manager or outside of NSX Manager. The **Service Certificate** flag is deactivated for CSRs generated on NSX Manager. Therefore, the signed CSRs cannot be used as service certificates, but only as platform certificates.

Platform and NSX service certificates are stored separately within the system and certificates imported as NSX service certificate cannot be used for platform or the reverse.

Principal Identity (PI) Certificates

API requests use PI certificates. PI certificates are one of the NSX Manager authentication mechanisms. Any NSX Manager client can create and use a PI certificate. It is the preferred approach for machine-to-machine communication.

PI for Cloud Management Platforms (CMP), such as Openstack, uses X.509 certificates that are uploaded when onboarding a CMP as a client. For information on assigning roles to Principal Identity and replacing PI certificates, see Add a Role Assignment or Principal Identity

PI for NSX Federation uses X.509 platform certificates for the Local Manager and Global Manager appliances. See Certificates for NSX and NSX Federation for details on self-signed certificates auto-configured for NSX Federation.

Certificates for NSX and NSX Federation

The system creates certificates required for the communication between NSX appliances and external communication, including NSX Federation appliances.

To replace existing self-signed certificates with CA-signed certificates, refer to details in Replace Certificates . To read about security compliance events, refer to the NSX Event Catalog.

Table 23-2. Certificates for NSX Managers

Certificate Naming Convention	Purpose	Replaceable using service_type	Default Validity
APH-AR	Appliance Proxy hub (APH) server public key and Aysnchronous Replicator	Yes, use service_type=APH.	825 days
APH-TN	Appliance Proxy hub (APH) certificate for Transport node (TN) and intra-cluster communication	Yes, use service_type=APH_TN.	825 days
API (aka tomcat)	API server certificate for NSX Manager node	Yes, use service_type=API.	825 days
API-Corfu Client	Corfu client certificate	Yes, use service_type=CBM_API.	100 years
AR-Corfu Client	Corfu client certificate	Yes, use service_type=CBM_AR.	100 years
CCP-Corfu Client	Control Configuration Plane Corfu client certificate	Yes, use service_type=CBM_CCP.	100 years
Cluster (aka mp-cluster)	Corfu client certificate	Yes, use service_type=MGMT_CLUSTER.	825 days
Cluster Manager-Corfu	Corfu client certificate	Yes, use service_type=CBM_CLUSTER_MANAGER.	100 years
CM Inventory-Corfu Client	Corfu client certificate	Yes, use service_type=CBM_CM_INVENTORY.	100 years
Corfu Server	Corfu client certificate	Yes, use service_type=CBM_CORFU.	100 years
IDPS reporting-Corfu Client	Corfu client certificate	Yes, use service_type=CBM_IDPS_REPORTING.	100 years
Messaging Manager-Corfu Client	Corfu client certificate	Yes, use service_type=CBM_MESSAGING_MANAGER.	100 years
Monitoring-Corfu Client	Corfu client certificate	Yes, use service_type=CBM_MONITORING.	100 years
MP-Corfu Client	Corfu client certificate	Yes, use service_type=CBM_MP.	100 years

Table 23-2. Certificates for NSX Managers (continued)

Certificate Naming Convention	Purpose	Replaceable using service_type	Default Validity
Site Manager-Corfu Client	Corfu client certificate	Yes, use service_type=CBM_SITE_MANAGER	100 years
Upgrade Coordinator-Corfu Client	Corfu client certificate	Yes, use service_type=CBM_UPGRADE_COORDINATOR.	100 years

Certificates for NSX Federation Communication

By default, the Global Manager uses self-signed certificates for communicating with internal components, registered Local Managers, and for authentication for NSX Manager UI or APIs.

You can view the external (UI/API) and inter-site certificates in NSX Manager. The internal certificates are not viewable or editable.

Note Do not enable Local Manager external VIP before you register the Local Manager on the Global Manager. When NSX Federation and PKS need to be used on the same Local Manager, complete the PKS tasks to create an external VIP and change the Local Manager certificate **before** you register the Local Manager on Global Manager.

By default, the Global Manager uses self-signed certificates for communicating with internal components, registered Local Managers, and for authentication for NSX Manager UI or APIs.

You can view the external (UI/API) and inter-site certificates in NSX Manager. The internal certificates are not viewable or editable.

Note Do not enable Local Manager external VIP before you register the Local Manager on the Global Manager. When NSX Federation and PKS need to be used on the same Local Manager, complete the PKS tasks to create an external VIP and change the Local Manager certificate **before** you register the Local Manager on Global Manager.

Certificates for Global Managers and Local Managers

After you add a Local Manager into the Global Manager, a trust gets established by exchanging certificates between Local Manager and Global Manager. These certificates also get copied into each of the sites registered with the Global Manager. Starting in NSX 4.1.0, the certificate used to establish trust with the Global Manager gets generated only when the Local Manager registers with the Global Manager. That same certificate gets deleted if the Local Manager moves out of the NSX Federation environment.

See the Certificates for Global Managers and Local Managers table for a list of all the NSX Federation specific certificates created for each appliance and the certificates these appliances exchange with each other:

Table 23-3. Certificates for Global Managers and Local Managers

Naming Convention in the Global Manager or Local Manager	Purpose	Replaceable?	Default Validity
The following are certificates specific to each NSX Federation **appliance.**			
`APH-AR certificate`	■ For the Global Manager and each Local Manager. ■ Used for inter-site communication using the AR channel (Async-Replicator channel).	Yes, use service_type=APH. See Replace Certificates .	10 years
`GlobalManager`	■ For the Global Manager. ■ PI certificate for theGlobal Manager.	Yes, use service_type=GLOBAL_MANAGER. Refer to Replace Certificates .	825 days
`Cluster certificate`	■ For the Global Manager and each Local Manager. ■ Used for UI/API communication with the VIP of the Global Manager or Local Manager cluster.	Yes, use service_type=MGMT_CLUSTER. Refer to Replace Certificates .	825 days
`API certificate`	■ For the Global Manager and each Local Manager. ■ Used for UI/API communication with individual Global Manager and Local Manager nodes for each of the locations added to the Global Manager.	Yes, use service_type=API. See Replace Certificates .	825 days
`LocalManager`	■ Starting with NSX 4.1, generates only when Local Manager servers are in the NSX Federation environment. The certificate gets deleted if the Local Manager moves out of the NSX Federation environment. ■ PI certificate for this specific Local Manager.	Yes, use service_type=LOCAL_MANAGER. See Replace Certificates .	825 days
The LM and the GM share their Cluster, API, and APH-AR certificates between them. If a certificate is CA-signed, the CA is synchronized, but that certificate is not.			

Principal Identity (PI) Users for NSX Federation

After you add a Local Manager to the Global Manager, the following PI users with corresponding roles get created.

Table 23-4. Principal Identity (PI) Users Created for NSX Federation

NSX Federation Appliance	PI User Name	PI User Role
Global Manager	`LocalManagerIdentity` One for each Local Manager registered with this Global Manager.	Auditor
Local Manager	`GlobalManagerIdentity`	Enterprise Admin

	`LocalManagerIdentity`
	One for each Local Manager registered with the same Global Manager. To get a list of all the Local Manager PI users because they are not visible in the UI, enter the following API: command
	`GET https://<local-mgr>/api/v1/trust-management/principal-identities`
		Auditor

Create a Certificate Signing Request File

Certificate signing request (CSR) is an encrypted text that contains specific information such as, organization name, common name, locality, and country. You send the CSR file to a certificate authority (CA) to apply for a digital identity certificate.

By default, the NSX CSR generation UI and API do not support the SAN field. To create a CSR with SAN, you can use an experimental API, /api/v1/trust-management/csrs-extended. For more information, see the *NSX API Guide*.

Prerequisites

To fill out the CSR file details, gather the information. You must know the FQDN of the server and the organizational unit, organization, city, state, and country.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Select **System > Certificates**.
3. Click the **CSRs** tab.
4. Click **Generate CSR** and select **Generate CSR** or **Generate CA CSR** from the dropdown menu.
5. Complete the file details.

Option	Description
Common Name	Enter the fully qualified domain name (FQDN) of your server. For example, test.vmware.com.
Name	Assign a name for your certificate.
Organization Unit	Enter the department in your organization that is handling this certificate For example, IT department.

Option	Description
Organization Name	Enter your organization name with applicable suffixes. For example, VMware Inc.
Locality	Add the city in which your organization is located. For example, Palo Alto.
State	Add the state in which your organization is located. For example, California.
Country/Region	Add your organization location. For example, United States (US).
Algorithm	Set the encryption algorithm for your certificate. ■ RSA encryption - used for digital signatures and encryption of the message. ■ ECDSA (Elliptic Curve Digital Signature Algorithm) encryption - used for EAL4+ compliance. The performance of this algorithm is more efficient than RSA algorithm.
Key Size	Set the key bits size of the encryption algorithm. ■ For RSA, the default value, 2048, is adequate unless you specifically need a different key size. Other supported sizes are 3072 and 4096. Many CAs require a minimum value of 2048. Larger key sizes are more secure but have a greater impact on performance. ■ ECDSA typically uses the Advanced Encryption Standard with 256 bit key in Galois/Counter mode (AES 256 GCM). Other key sizes include 384 and 521 bits.
Description	Enter specific details to help you identify this certificate at a later date.

6 Click **Save**.

A custom CSR appears as a link.

7 Select the CSR then click **Actions** to select one of the following options:

- **Delete**
- **Import Certificate for CSR**
- **Self Sign Certificate for CSR**
- **Download CSR PEM**

  If you selected **Download CSR PEM**, you can save the CSR PEM file for your records and CA submission. Use the contents of the CSR file to submit a certificate request to the CA in accordance with the CA enrollment process. For the other two options, refer to topics Import a Certificate for a CSR and Create a Self-Signed Certificate.

Results

The CA creates a server certificate based on the information in the CSR file, signs it with its private key, and sends you the certificate. The CA also sends you a root CA certificate.

Creating Self-signed Certificates

You can create self-signed service or non-service certificates in NSX Manager. You can also import a signed certificate for NSX-generated CSRs.

Create a Self-Signed Certificate

You can create a self-signed service or non-service certificate. However, using a self-signed certificate is less secure than using a trusted certificate.

When you use a self-signed certificate the client user receives a warning message such as, `Invalid Security Certificate`. The client user must then accept the self-signed certificate when first connecting to the server in order to proceed. Allowing client users to select this option provides reduced security than other authorization methods.

Prerequisites

Verify that a CSR is available. See Create a Certificate Signing Request File.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Select **System > Certificates**.
3. Click the **CSRs** tab.
4. From your selected CSR, click  and select **Self Sign Certificate for CSR**.

   Note If you have a self signed CA CSR, NSX Manager always creates a CA CSR.
5. Enter the number of days the self-signed certificate is valid.

   The default is 825 days. Even if you change this value for previously generated self-signed certificate, the default value is displayed every time you generate a new certificate.
6. Choose your **Service Certificate** type.
   a. Toggle the **Service Certificate** button to **Yes** to use this certificate for services such as load balancer, VPN, or TLS Inspection. If you are creating a self-signed CA certificate, Yes is the only choice.
   b. Toggle the **Service Certificate** button to **No** to use this certificate with NSX Manager appliance nodes.
7. Click **Save**.

Results

The self-signed certificate appears in the **Certificates** tab.

Import a Certificate for a CSR

You can import a signed certificate for an NSX generated CSR (certificate signing request). You can also use this imported certificate with services such as Load Balancer, VPN, and TLS Inspection. This page provides the steps to import a signed certificate for NSX generated CSR.

A self-signed certificate acts as a certificate as well as CA. It is not required to be signed from any external CA, whereas CSR is a certificate signing request that cannot act as CA and must be signed by external CA. There is no support for a self-signed certificate for load balancer.

When you use a self-signed certificate the client user receives a warning message such as, `Invalid Security Certificate`. The client user must then accept the self-signed certificate when first connecting to the server in order to proceed. Allowing client users to select this option provides reduced security than other authorization methods.

Prerequisites

- Verify that a CSR is available. See Create a Certificate Signing Request File.
- NSX generated CSR was used as a CSR for signed certificate.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Select **System > Certificates**.
3. Click the **CSRs** tab.
4. From a CSR, click  and select **Import Certificate for CSR**.
5. Browse to the signed certificate file on your computer and add the file.
6. Choose your **Service Certificate** type.
   a. To use this certificate for services such as load balancer, VPN, or TLS Inspection, toggle the **Service Certificate** button to **Yes**.
   b. To use this certificate with NSX Manager appliance nodes, toggle the **Service Certificate** button to **No**.
7. Click **Save**.

Results

The signed certificate appears in the **Certificates** tab.

Importing and Replacing Certificates

You can import self-signed or CA-signed certificates for platform or services. You can replace the self-signed certificates using APIs.

You can also import CA certificates for services such as Load Balancer.

Import a Self-signed or CA-signed Certificate

You can import a certificate with a private key to replace the default self-signed certificate, after activation.

You can import self-signed or CA-signed certificates for platform or services using this procedure. Note that a CSR generated on NSX Manager and which is self-signed cannot be used as a service certificate, such as the Load Balancer service. If you want to import a CA certificate for the Load Balancer service, see Import a CA Certificate.

Prerequisites

- Verify that a certificate is available.
- The server certificate must contain the Basic Constraints extension `basicConstraints = cA:FALSE`.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Select **System > Certificates**.
3. Select **Import > Import Certificate** and enter the certificate details.

Option	Description
Name	Assign a name to the certificate.
Service Certificate	Set to **Yes** to use this certificate for services such as a load balancer and VPN. Set to **No** if this certificate is for the NSX Manager nodes.
Certificate Contents	Browse to the certificate file on your computer and add the file. The certificate must not be encrypted. If it is a CA-signed certificate, be sure to include the whole chain in this order: certificate - intermediate - root.
Private Key	Browse to the private key file on your computer and add the file. Private key is an optional field if imported certificate is based on NSX Manager generated CSR, as a private key exists on the NSX Manager appliance.
Passphrase	In this release, this field is not used.
Description	Enter a description of what is included in this certificate.

4. Click **Import**.

Import a CA Certificate

You can import a CA certificate from a system external to NSX, for example, to use with the Load Balancer service.

Prerequisites

Verify that a CA certificate is available.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **System > Certificates**.

3 Select **Import > CA Certificate** and enter the certificate details.

Option	Description
Name	Assign a name to the CA certificate.
Service Certificate	Set **Yes** to use this certificate for services such as a load balancer, VPN, or TLS Inspection. Set to **No** to use the certificate with NSX Manager appliance nodes.
Certificate Contents	Browse to the CA certificate file on your computer and add the file.
Private key	For service certificate only. Browse to the private key file on your computer and add the file. Private key is an optional field. Sent to the edge node where the feature is running.
Passphrase	In this release, this field is not used.
Description	Enter a description of what is included in this certificate.

4 Click **Save**.

Set Checks for Certificate Imports

You can enable or disable the Extended Key Usage (EKU) Extension and the Certificate Revocation List Distribution Point (CDP) validation checks that NSX performs while importing a certificate.

Note: If you have CA-signed certificates without a CDP then you might have problems after upgrade. To avoid this problem you can turn CRL checking off or replace the certificates with certificates that include a CDP.

To set validation checks, use the following API with payload. For more information about the API, see the *NSX API Guide*.

```
PUT https://<manager>/api/v1/global-configs/SecurityGlobalConfig
{
"crl_checking_enabled": false,
"ca_signed_only": false,
"eku_checking_enabled":false,
"resource_type":"SecurityGlobalConfig"
}
```

Where:

- `crl_checking_enabled`: Enabled by default to check CDP specified in the imported CA-signed certificate. Support includes HTTP based CRL-DP only. File or LDAP-based options are not supported.
- `ca_signed_only`: Disabled by default. It allows checks signed by CA only.

- `eku_checking_enabled`: Disabled by default. It checks for EKU Extension in the imported certificate.
- `revision`: The current revision of the resource that must be included in a request. To obtain the value of this parameter issue a GET operation.

Replace Certificates

After you install NSX the manager nodes and cluster have self-signed certificates. Replace the self-signed certificates with CA-signed certificates and use a single common CA-signed certificate with a SAN (Subject Alternative Names) list that matches all the nodes and the VIP for the cluster.

If you are using NSX Federation, you can replace the GM API certificates, GM cluster certificate, LM API certificates, and LM cluster certificates using the following APIs.

Started in NSX Federation 4.1, you can replace the self-signed certificate used for the GM-LM communication. Also, the Global Manager certificate now generates the Local Manager certificate at the time the Local Manager gets registered. The Local Manager certificate is no longer a default certificate.

When you replace the Global Manager or Local Manager certificate, the site-manager sends these to all the other federated sites, so communication should remain intact.

The cipher suite TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384 can now be used or replaced for communication between:

- the NSX nodes with in the cluster
- within the NSX Federation
- NSX Manager to NSX Edge
- NSX Manager to NSX agent
- NSX Manager REST API communication (external)

You can also replace the platform Principal Identity certificates auto-created for the Global Manager and Local Manager appliances. See Certificates for NSX and NSX Federation for details on self-signed certificates auto-configured for NSX Federation.

Note For Cloud Service Manager, it is not possible to replace the HTTP certificate in an NSX environment.

Prerequisites

- Verify that a certificate is available in the NSX Manager. See Import a Self-signed or CA-signed Certificate.
- The server certificate must contain the Basic Constraints extension `basicConstraints = CA:FALSE`.
- Verify that the certificate is valid by making the following API call:

```
GET https://<nsx-mgr>/api/v1/trust-management/certificates/<cert-id>?
action=validate
```

-

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **System > Certificates**.

3 In the ID column, select the ID of the certificate you want to use and copy the certificate ID from the pop-up window.

Make sure that when this certificate was imported, the option **Service Certificate** was set to **No**.

Note: The certificate chain must be in the industry standard order of certificate - intermediate - root.

4 To replace the API certificate of a manager node, use the API call:

```
POST /api/v1/trust-management/certificates/<cert-id>?
action=apply_certificate&service_type=API&node_id=<node-id>
```

For example:

```
POST
https://<nsx-mgr>/api/v1/trust-management/certificates/77c5dc5c-6ba5-4e74-a801-
c27dc09be76b?action=apply_certificate&service_type=API&node_id=e61c7537-3090-4149-
b2b6-19915c20504f
```

For more information about the API, see the *NSX API Guide*.

5 To replace the certificate of the manager cluster VIP, use the API call:

```
POST /api/v1/trust-management/certificates/<cert-id>?
action=apply_certificate&service_type=MGMT_CLUSTER
```

For example:

```
POST https://
<nsx-mgr>/api/v1/trust-management/certificates/d60c6a07-6e59-4873-8edb-339bf75711?
action=apply_certificate&service_type=MGMT_CLUSTER
```

For more information about the API, see the *NSX API Guide*. This step is not necessary if you did not configure the VIP.

6 (Optional) To replace the Local Manager and Global Manager Principal Identity certificates for NSX Federation use the API call. The entire NSX Manager cluster (Local Manager and Global Manager) requires a single PI certificate.

```
POST https://<nsx-mgr>/api/v1/trust-management/certificates/<cert-id>?
action=apply_certificate&service_type=<service-type>
```

For example:

```
POST https://<local-mgr>/api/v1/trust-management/certificates/77c5dc5c-6ba5-4e74-a801-
c27dc09be76b?action=apply_certificate&service_type=LOCAL_MANAGER
```

or

```
POST https://<global-mgr>/api/v1/trust-management/certificates/77c5dc5c-6ba5-4e74-a801-
c27dc09be76b?action=apply_certificate&service_type=GLOBAL_MANAGER
```

7 To replace APH-AR certificates use the API call:

```
POST https://<nsx-mgr>/api/v1/trust-management/certificates/<cert-id>?
action=apply_certificate&service_type=APH
```

For example:

```
POST https://<nsx-mgr>/api/v1/trust-management/certificates/77c5dc5c-6ba5-4e74-a801-
c27dc09be76b?action=apply_certificate&service_type=APH
```

8 Started in NSX 4.1, to replace the transport node (TN) or the Edge host certificates use the API call:

```
POST https://<nsx-mgr>/api/v1/trust-management/certificates/action/<cert-id>?
action=replace-host-certificate/<host-id or node-id>
{
      "pem_encoded": "---BEGIN CERTIFICATE---<certificate>---END CERTIFICATE-----\n-----
BEGIN CERTIFICATE---<certificate>---END CERTIFICATE---\n",
      "private_key": "---BEGIN RSA PRIVATE KEY---<private rsa key>---END RSA PRIVATE KEY---
\n"
    }
```

For example:

```
POST https://<nsx-mgr>/api/v1/trust-management/certificates/action= replace-host-
certificate/host-id=0e83d0124-75b1-d6ee-b37c-ce8f22d71204
{
      "pem_encoded": "---BEGIN CERTIFICATE---C1tZ3…Jlmno338p---END CERTIFICATE---\n---
BEGIN CERTIFICATE---d9f6jj01…P241kwm06---END CERTIFICATE---\n",
      "private_key": "---BEGIN RSA PRIVATE KEY---IB4AA5K3…C8A12Q7E---END RSA PRIVATE KEY---
\n"
}
```

Use the ID of the transport node for the node_id, not the Unified Appliance (UA) ID.

For details on the default self-signed certificates configured by the system or for VMware Cloud Foundation certificate requirements, refer to Types of Certificates.

Importing and Retrieving CRLs

You can import a certificate revocation list (CRL) into the NSX Manager and use the API to configure NSX Manager to retrieve a CRL.

Import a Certificate Revocation List

A certificate revocation list (CRL) is a list of subscribers and their certificate status. When a potential user attempts to access a server, the server denies access based on the CRL entry for that particular user. This topic describes how to import a CRL into the NSX Manager.

NSX supports two CRL formats:

- PEM-encoded X.509 CRL - 40 MB maximum size, 500,000 entries
- Mozilla OneCRL - 5 MB maximum size, 10,000 entries

The list contains the following items:

- Revoked certificates and the reasons for revocation
- Dates the certificates are issued
- Entities that issued the certificates
- Proposed date for the next release

Prerequisites

Verify that a CRL is available.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Select **System > Certificates**.
3. Click the **CRLs** tab.
4. To browse the *default_public_crl* file, expand that row and click **View Details**.

   You can view the Issuer Name and Serial Numbers details.
5. To import a CRL, click **Import** and add the CRL details.

Option	Description
Name	Assign a name to the CRL.
CRL Bundle	Browse for your PEM or JSON files and select the file for import.
Description	Enter a summary of what is included in this CRL.

6. Click **Save**.

Results

The imported CRL appears as a link.

Configuring NSX Manager to Retrieve a Certificate Revocation List

Using the API, you can configure NSX Manager to retrieve a certificate revocation list (CRL). You can then check the CRL by making an API call to NSX Manager instead of to the certificate authority.

This feature provides the following benefits:

- It is more efficient to have the CRL cached on the server, that is, NSX Manager.
- The client does not need to create any outbound connection to the certificate authority.

The following APIs related to certificate revocation lists are available:

```
GET /api/v1/trust-management
GET /api/v1/trust-management/crl-distribution-points
POST /api/v1/trust-management/crl-distribution-points
DELETE /api/v1/trust-management/crl-distribution-points/<crl-distribution-point-id>
GET /api/v1/trust-management/crl-distribution-points/<crl-distribution-point-id>
PUT /api/v1/trust-management/crl-distribution-points/<crl-distribution-point-id>
GET /api/v1/trust-management/crl-distribution-points/<crl-distribution-point-id>/status
POST /api/v1/trust-management/crl-distribution-points/pem-file
```

You can manage CRL distribution points and retrieve the CRLs stored in NSX Manager. For more information, see the *NSX API Guide*.

Import or Update a Trusted CA Bundle

You can now use a built-in trusted certificate authority (CA) bundle for the TLS Inspection chain of trust to support advanced security applications such as IDS/IPS, URL filtering, malware, and granular App ID.

You can use the built-in CA bundle, `default_trusted_public_ca_bundle`, internally for the TLS inspection and decryption for gateway firewalls.

For external services, TLS Proxy requires a configured trusted CA bundle to validate the certificate that any external service presents to it. You can configure the External_Decryption_Profile.trusted_ca_bundles with one or more CA bundles where each bundle is a list of certificates. You must configure at least one CA bundle. Typically, external services use well known CAs such as Verisign and DigiCert. So, for ease of configuration, NSX includes a built-in default_trusted_public_ca_bundle that contains a list of widely used CA certs, similar to how operating systems come pre-installed with popular CA certs. You can update this bundle or you can create your own CA bundle and use it instead.

You can perform the following tasks in NSX. You can find Trusted CA Bundles by selecting **System > Certificates > Trusted CA Bundle**.

- Validate TLS inspection and decryption using the default trusted CA bundle.
- View all certificates in the CA bundle including filtering basic details using the **View All Certificates** button.

- Search for expired, expiring, valid, used and unused CA bundles using the **View All Certificates** button.
- Edit CA bundle display name and add or remove certificates from the bundle.
- Export a CA bundle for inclusion on other devices.
- Copy the CA bundle path locally.
- Import a new trusted CA bundle using the **Import CA Bundle** button.

Storage of Public Certificates and Private Keys for Load Balancer or VPN service

Public certificates and private keys are stored on the NSX Managers for load balancer or VPN service. When a load balancer or VPN service is created that requires a private key, NSX Manager sends a copy of the private key to the Edge node where the load balancer or VPN service is running.

Alarm Notification for Certificate Expiration

NSX generates alarms when a certificate is nearing its expiry or if a certificate has already expired. Service certificates generate an alarm only if expiring or expired and in use by a component. Non-service certificates always generate an alarm, whether in use or not.

NSX generates alarms under following events. The defaults are listed below, but are configurable.

- Medium severity alarm starting 30 day before certificate expiry.
- High severity alarm starting 7 days prior to expiry.
- Critical severity alarm every day after certificate expires.

Certificate Expiry alarms contains details on certificate ID, severity, node, first/last report time, and recommended action.

As a remedial, you must replace the expiring External Platform certificate with a new valid certificate and delete expiring certificate.

Integration of Kubernetes Clusters with Antrea CNI

24

Antrea is a container network interface (CNI) plugin from VMware that provides network connectivity and security features to pods in container clusters that are based on Kubernetes.

Note Henceforth, this documentation uses the term "Antrea Kubernetes cluster" to mean Kubernetes clusters with Antrea CNI. The term "Kubernetes cluster" is a generic term, which represents Tanzu Kubernetes Grid (TKG) clusters with Antrea CNI, OpenShift clusters with Antrea CNI, or do it yourself (DIY) Kubernetes clusters with Antrea CNI.

The objective is to connect Antrea Kubernetes clusters to the NSX Management Plane and Central Control Plane (CCP). To achieve this integration, you must deploy Antrea NSX Adapter on all the Antrea Kubernetes clusters that you want to integrate to NSX.

Benefits of Integration

The integration of Antrea Kubernetes clusters to NSX enables the following capabilities:

- View Antrea Kubernetes cluster resources in the NSX Manager UI (Policy mode).
- Centrally manage groups and security policies in NSX that reference Antrea Kubernetes clusters and NSX resources (for example, VMs).
- Distribute the NSX security policies to the Kubernetes clusters for enforcement in the cluster by the Antrea CNI.
- Extend the NSX network diagnostic and troubleshooting features to the Antrea Kubernetes clusters, such as collecting support bundles, monitoring logs, and performing Traceflow operations.
- Monitor the runtime state and health status of Antrea Kubernetes cluster components and Antrea Agents in the NSX Manager UI.

Interoperability Requirements

For NSX and Antrea integration, specific interoperability requirements must be met. For more details, see VMware Product Interoperability Matrix.

This chapter includes the following topics:

- Architecture of Antrea Kubernetes Cluster Integration with NSX

- Registering an Antrea Kubernetes Cluster to NSX
- Viewing Inventory of an Antrea Kubernetes Cluster in NSX Manager
- Monitor Health Status of an Antrea Kubernetes Cluster
- Trace the Path of a Packet with Antrea Traceflow
- Antrea Groups
- Add an Antrea Group
- Generic Groups with Kubernetes Member Types in Dynamic Membership Criteria
- Distributed Firewall Policies for Securing Traffic Within an Antrea Kubernetes Cluster
- Firewall Policies for Securing Traffic Between Antrea Kubernetes Clusters and VMs in an NSX Network
- Deregister an Antrea Kubernetes Cluster from NSX
- Troubleshooting Antrea to NSX Integration Issues

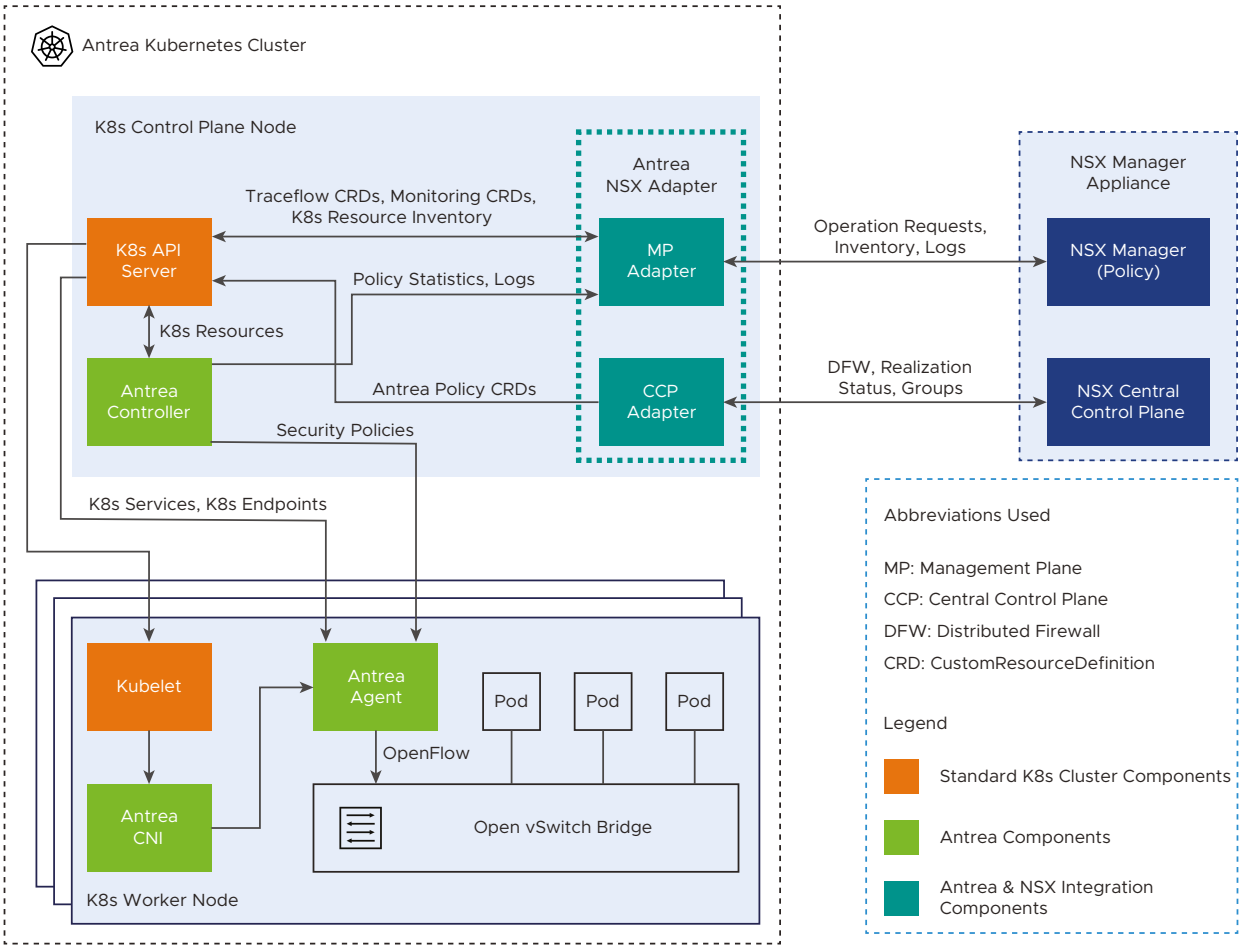
Architecture of Antrea Kubernetes Cluster Integration with NSX

The integration architecture explains the information exchanged between a Kubernetes cluster that uses Antrea CNI and the NSX Manager Appliance, which is deployed in NSX.

This documentation does not explain the functions of Antrea components in a Kubernetes (K8s) cluster. To understand the Antrea architecture and the functions of Antrea components in a Kubernetes cluster, see the Antrea documentation portal at https://antrea.io/docs.

This main objective of this documentation is to understand the functions of the Antrea NSX Adapter that integrates a Kubernetes cluster with Antrea CNI to the NSX Manager Appliance.

Architecture Diagram



Antrea NSX Adapter

This component runs as a pod on one of the Kubernetes Control Plane nodes. Antrea NSX Adapter consists of the following two subcomponents:

- Management Plane Adapter (MP Adapter)
- Central Control Plane Adapter (CCP Adapter)

Management Plane Adapter communicates with the NSX Management Plane (Policy), Kubernetes API Server, and Antrea Controller. Central Control Plane Adapter communicates with the NSX Central Control Plane (CCP) and Kubernetes API Server.

Functions of the Management Plane Adapter

- Watches the Kubernetes resource inventory from Kubernetes API and reports the inventory to NSX Manager. Resource inventory of an Antrea Kubernetes cluster includes resources, such as Pods, Ingress, Services, Network Policies, Namespaces, and Nodes.

- Responds to the policy statistics query from NSX Manager. It receives the statistics from the Antrea Controller API or the statistics that are exported by the Antrea Agent on each K8s worker node, and reports the statistics to NSX Manager.
- Receives troubleshooting operation requests from NSX Manager, sends the requests to Antrea Controller API server, collects the results, and returns the information to NSX Manager. Examples of troubleshooting operations include Traceflow requests, Support Bundle collection requests, log collection requests.
- Watches the runtime state and health status of an Antrea Kubernetes cluster from the Antrea Monitoring CustomResourceDefinition (CRD) objects and reports the status to NSX Manager. The status is reported on a per cluster basis. For example, the health status of the following components is reported to NSX Manager:
  - Management Plane Adapter
  - Central Control Plane Adapter
  - Antrea Controller
  - Antrea Agents

Functions of the Central Control Plane Adapter

- Receives the Distributed Firewall (DFW) rules and groups from NSX Central Control Plane, translates them to Antrea policies, and creates Antrea policy CRDs in K8s API.
- Watches the policy realization status from both K8s network polices and native Antrea policy CRDs and reports the status to NSX Central Control Plane.

Stateless Nature of the Central Control Plane Adapter

The Central Control Plane Adapter is stateless. Each time the adapter restarts or reconnects to K8s API or NSX Manager, it always synchronizes the state with K8s API and NSX Central Control Plane. Resynchronization of the state ensures the following:

- The latest Antrea policies are always pushed to K8s API as native Antrea policy CRDs.
- The stale policy CRDs are removed if the corresponding security policies are deleted in NSX.

Registering an Antrea Kubernetes Cluster to NSX

The registration process involves two roles or personas: NSX administrator and Kubernetes platform administrator. In an organization, both roles might be held by the same person or by different people.

For example, if your organization uses a hosted (managed) Kubernetes service from a public cloud service provider, such as Azure Kubernetes Service, Google Kubernetes Service, and so on, the Kubernetes platform administrator is in the cloud service provider organization, and the NSX administrator is in your organization. In such a scenario, the Kubernetes platform administrator and the NSX administrator can collaborate with each other to exchange information and complete the registration.

Prerequisites for Registering an Antrea Kubernetes Cluster to NSX

Before registering an Antrea Kubernetes cluster to NSX, you must complete several prerequisite tasks.

You can register multiple Antrea Kubernetes clusters to a single NSX deployment.

Ensure that Required Ports are Opened for Antrea-NSX Interworking Adapter

Antrea-NSX Interworking Adapter runs as a pod in an Antrea Kubernetes cluster, and this pod uses the host networking mode. The pod can be scheduled to run on any Kubernetes node. Therefore, you must ensure that the K8s nodes can reach the NSX IP addresses on the ports that are mentioned on the VMware Ports and Protocols portal at https://ports.esp.vmware.com/home/NSX.

At this link, enter **`Antrea Interworking Pod`** in the **Search** text box.

Deploy Antrea Kubernetes Clusters

Persona: Kubernetes platform administrator

A Kubernetes cluster with Antrea network plug-in must be up and ready.

For example, to integrate clusters in a Tanzu Kubernetes Grid instance with NSX, ensure that the following tasks are completed:

- Tanzu management clusters are deployed and the clusters are in running state.
- Tanzu Kubernetes clusters are deployed and the clusters are in running state.
- Tanzu command line interface (CLI) is installed.

For a detailed information about these tasks, see the *Tanzu Kubernetes Grid* documentation at https://docs.vmware.com/en/VMware-Tanzu-Kubernetes-Grid/index.html.

When you deploy a management cluster, networking with Antrea is automatically enabled in the management cluster.

Add an Appropriate License in NSX

Persona: NSX administrator

Ensure that your NSX environment has one of these licenses:

- NSX Data Center Advanced
- NSX Data Center Enterprise Plus

- Antrea Enterprise Standalone

To add a license:

1 In NSX Manager, navigate to **System > Licenses > Add Licenses**.

2 Enter a license key.

Create a Self-Signed Security Certificate

Persona: NSX administrator

A self-signed security certificate is required to create a principal identity user account in NSX, which is explained later in this topic.

Using OpenSSL commands, create a self-signed security certificate for each Antrea Kubernetes cluster that you want register to NSX.

For example, assume that you want to create a self-signed OpenSSL certificate of length 2048 bits for an Antrea Kubernetes cluster called cluster-sales. The following OpenSSL commands generate a private key file, a certificate signing request file, and a self-signed certificate file for this cluster.

```
openssl genrsa -out cluster-sales-private.key 2048
openssl req -new -key cluster-sales-private.key -out cluster-sales.csr -subj "/C=US/ST=CA/
L=Palo Alto/O=VMware/OU=Antrea Cluster/CN=cluster-sales"
openssl x509 -req -days 3650 -sha256 -in cluster-sales.csr -signkey cluster-sales-private.key
-out cluster-sales.crt
```

Note In the `openssl req` command that you use to create the `.csr` file, ensure that the Common Name (CN) is different for each Antrea Kubernetes cluster.

Create a Principal Identity User

Persona: NSX administrator

The Management Plane Adapter and Central Control Plane Adapter use the principal identity (PI) user account to authenticate with an NSX Manager and identify themselves as the principal identity. The PI user owns the inventory resources that are reported by the adapters. NSX prevents other users from accidentally overwriting the inventory resources.

We recommend that you use a separate PI user account for registering each Antrea Kubernetes cluster so that identification of objects becomes easier. However, it is fine to use an existing PI user account. Creating a new PI user account for each Antrea Kubernetes cluster is not mandatory.

Create a principal identity user in NSX with the self-signed certificate that you created in the previous step. Assign this principal identity user an **Enterprise Admin** role.

To create a principal identity user:

1 In the NSX Manager UI, click the **System** tab.

2 Under **Settings**, navigate to **User Management > User Role Assignment**.

3 Click **Add > Principal Identity with Role**.

4 Enter a name for the principal identity user. For example, enter **`cluster-sales`**.

Important Ensure that you specify the same name for the NSX principal identity user, certificate CN, and the clusterName argument in the `bootstrap-config.yaml` file.

For more information about the bootstrap configuration file, see Edit the Bootstrap Configuration File.

5 Select the role as **Enterprise Admin**.

6 In the **Node Id** text box, enter a name for the Antrea Kubernetes cluster. This name must be unique across all Kubernetes clusters that you are registering to NSX. For example, enter **`cluster-sales`**.

7 In the **Certificate PEM** text area, paste the complete self-signed certificate, which you created earlier. Ensure that the -----BEGIN CERTIFICATE---- and ------END CERTIFICATE----- lines are also pasted in this text box.

8 Click **Save**.

9 From the left navigation pane, under **Settings**, click **Certificates**. Verify that the self-signed certificate of the Antrea Kubernetes cluster is shown.

Determine the Antrea Version From the Kubernetes Cluster

Persona: Kubernetes platform administrator

Before downloading the Antrea-NSX interworking file (`antrea-interworking-`*`version`*`.zip`), which is the next prerequisite in this topic, you must determine the Antrea open source version from your Kubernetes cluster.

Important Each VMware Container Networking™ with Antrea™ release is based on one Antrea open source version. Antrea-NSX interworking version is compatible with the Antrea open source software version from the same VMware Container Networking release.

For example, see the following table.

VMware Container Networking Version	Based on Antrea OSS Version	Compatible With Antrea-NSX Interworking Version
v1.5.0 See: v1.5.0 Release Notes	v1.7.1	v0.7.*
v1.4.0 See: v1.4.0 Release Notes	v1.5.2	v0.5.*
v1.3.1 See: v1.3.1-1.2.3 Release Notes	v1.2.3	v0.2.*

To determine the Antrea open source version from your Kubernetes cluster, do these steps:

1 Find out the Antrea Controller pod name. Kubernetes generates this name with a random string, so, you can get the name from the K8s cluster.

For example:

```
$ kubectl get pod -n kube-system -l component=antrea-controller
NAME                                READY   STATUS    RESTARTS   AGE
antrea-controller-6b8cb7cd59-wcjvd   1/1     Running   0          13d
```

In this command output, the Antrea Controller pod name is `antrea-controller-6b8cb7cd59-wcjvd`.

2 Retrieve the Antrea open source version by running the following command:

```
$ kubectl exec -it antrea-controller-6b8cb7cd59-wcjvd -n kube-system -- antctl version
antctlVersion: v1.7.1-cacafc0
controllerVersion: v1.7.1-cacafc0
```

In this command output, `v1.7.1` is the Antrea open source version that you wanted to determine.

Download the Antrea-NSX Interworking Zip File

Persona: Kubernetes platform administrator

Complete the following steps to download the `antrea-interworking-`*`version`*`.zip` file:

1 In a Web browser, open the **All Downloads** page on the VMware Customer Connect portal, and log in with your VMware ID.

2 From the **All Products** drop-down menu, select **Networking & Security**.

3 Next to **VMware Antrea**, click **Download Product**. The **Download VMware Antrea** page opens.

4 Next to the **VMware Container Networking with Antrea** version that is relevant to you, click **Go to Downloads**.

5 Find **Antrea-NSX Interworking Adapter Image and Manifest Bundle** on the page. Click **Read More** and verify that the version mentioned in the `antrea-interworking-`*`version`*`.zip` file is compatible with the Antrea open source software version, which you determined earlier.

6 Click **Download Now**.

Extract the ZIP file. It contains the following files.

File Name	Description
`interworking.yaml`	YAML deployment manifest file to register an Antrea Kubernetes cluster to NSX.
`bootstrap-config.yaml`	YAML file where you can specify the following details for registration: Antrea Kubernetes cluster name, NSX Manager IP addresses, TLS certificate of the Kubernetes cluster, and the private key of the Kubernetes cluster.
`deregisterjob.yaml`	YAML manifest file to deregister an Antrea Kubernetes cluster from NSX.
`ns-label-webhook.yaml`	Webhook definitions for automatically adding labels to newly created Kubernetes namespaces. This YAML file is used only when Kubernetes version is ≤ 1.20.
`interworking-`*`version`*`.tar`	Archive file for the container images of Management Plane Adapter and Central Control Plane Adapter.

Import the Container Images to Container Registry

Persona: Kubernetes platform administrator

There are two approaches for doing this prerequisite task.

Approach 1 (Recommended): Pull images from VMware Harbor Registry

VMware has hosted the container images on VMware Harbor Registry.

Image locations are as follows:

- `projects.registry.vmware.com/antreainterworking/interworking-debian:`*`version`*
- `projects.registry.vmware.com/antreainterworking/interworking-ubuntu:`*`version`*
- `projects.registry.vmware.com/antreainterworking/interworking-photon:`*`version`*

For *version* information, see the *VMware Container Networking with Antrea* release notes at https://docs.vmware.com/en/VMware-Container-Networking-with-Antrea/index.html.

Open the `interworking.yaml` and `deregisterjob.yaml` files in any text editor of your choice, and replace all image URLs with any one of these image locations.

The advantage of this approach is that when you submit the `.yaml` files to the Kubernetes API server for registering the cluster, Kubernetes can pull the container images automatically from VMware Harbor Registry.

Approach 2: Manually copy images to Kubernetes worker nodes and control plane nodes

If your Kubernetes infrastructure has no Internet connectivity, or connectivity is too slow, use this manual approach.

Extract the container images from the `interworking-version.tar` file and copy them to the Kubernetes worker nodes and control plane node of each Antrea Kubernetes cluster that you want to register to NSX.

For example, at the Tanzu CLI, run the following command for each Kubernetes worker node IP and control plane node IP to copy the `.tar` and `.yaml` files:

```
scp -o StrictHostKeyChecking=no interworking* capv@{node-ip}:/home/capv
```

Import the images to the local Kubernetes registry, which is managed by the container runtime engine. Alternatively, if your organization has a private container registry, you can import the container images to the private container registry.

For example, at the Tanzu CLI, run the following command for each Kubernetes worker node IP and control plane node IP to import the container images to the local Kubernetes registry:

```
ssh capv@{node-ip} sudo ctr -n=k8s.io i import interworking-{version-id}.tar
```

Edit the Bootstrap Configuration File

In the `bootstrap-config.yaml` file, enter values of mandatory arguments that are required to register an Antrea Kubernetes cluster to NSX.

Prerequisites

Ensure that you complete the prerequisite tasks for registering an Antrea Kubernetes cluster to NSX. See Prerequisites for Registering an Antrea Kubernetes Cluster to NSX.

Bootstrap Configuration File

When you extract the `antrea-interworking-version.zip` file, you get a `bootstrap-config.yaml` file, which has some placeholder comments to help you fill the arguments in this file.

```
apiVersion: v1
kind: Namespace
metadata:
  name: vmware-system-antrea
  labels:
    app: antrea-interworking
    openshift.io/run-level: '0'
---
# NOTE: In production the bootstrap config and secret should be filled by admin
# manually or external automation mechanism.
apiVersion: v1
kind: ConfigMap
metadata:
  name: bootstrap-config
  namespace: vmware-system-antrea
```

```
data:
  bootstrap.conf: |
    # Fill in the cluster name. It should be unique across all clusters managed by NSX.
    clusterName: Name
    # Fill in the NSX manager IPs. If there are multiple IPs, separate them with commas.
    NSXManagers: [IP1, IP2, IP3]
    # vhcPath is deprecated and ignored. It is replaced by vpcPath.
    vhcPath: ""
    # vpcPath is supported only in Antrea-NSX interworking image v0.7 or later.
    vpcPath: ""
---
apiVersion: v1
kind: Secret
metadata:
  name: nsx-cert
  namespace: vmware-system-antrea
type: kubernetes.io/tls
data:
  # One line base64 encoded data. Can be generated by command: cat tls.crt | base64 -w 0
  tls.crt:
  # One line base64 encoded data. Can be generated by command: cat tls.key | base64 -w 0
  tls.key:
```

Specify Values of Mandatory Arguments

Each Kubernetes cluster that you want to register to NSX requires a separate `bootstrap-config.yaml` file. Specify values for the following mandatory arguments in this file.

clusterName

Enter a unique name for the Antrea Kubernetes cluster. This name must be unique across all Kubernetes clusters that are registered to NSX. For example, **`cluster-sales`**.

NSXManagers

Enter an NSX Manager IP address. To specify IP addresses of multiple NSX Manager nodes in an NSX Manager cluster, separate the IP addresses with a comma.

Note You can register an Antrea Kubernetes cluster to a single NSX Manager cluster. A single cluster can have one to three NSX Manager nodes.

For example, if your NSX Manager cluster has three NSX Manager nodes, enter the configuration as **`[192.168.1.1, 192.168.1.2, 192.168.1.3]`**. If your cluster has a single NSX Manager node, enter the configuration as **`[192.168.1.1]`**.

tls.crt

`tls.crt` is the self-signed certificate that you used to create the principal identity user in NSX.

Enter the one-line base64 encoded data of the TLS certificate for your Kubernetes cluster.

For example, to print the encoded data of the `cluster-sales.crt` certificate file on your terminal, run the following command at the Linux CLI:

```
cat cluster-sales.crt | base64 -w 0
```

tls.key

Enter the one-line base64 encoded data of the private key file for your Kubernetes cluster.

For example, to print the encoded data of the `cluster-sales-private.key` file on your terminal, run the following command at the Linux CLI:

```
cat cluster-sales-private.key | base64 -w 0
```

`tls.key` is not sent to NSX. The Management Plane Adapter and Central Control Plane Adapter use this key to prove that it owns the principal identity user.

Note The `vhcPath` argument is deprecated and ignored, whereas the `vpcPath` argument is supported. However, the `vpcPath` argument is supported only in Antrea-NSX interworking image v0.7 or later, which is included with VMware Container Networking™ with Antrea™ v1.5.0. Earlier versions of Antrea-NSX interworking image does not support the `vpcPath` argument.

Example: Bootstrap Configuration File

```
apiVersion: v1
kind: Namespace
metadata:
  name: vmware-system-antrea
  labels:
    app: antrea-interworking
    openshift.io/run-level: '0'
---
# NOTE: In production the bootstrap config and secret should be filled by admin
# manually or external automation mechanism.
apiVersion: v1
kind: ConfigMap
metadata:
  name: bootstrap-config
  namespace: vmware-system-antrea
data:
  bootstrap.conf: |
    clusterName: cluster-sales
    NSXManagers: [10.196.239.128, 10.196.239.129]
    vpcPath: ""
---
apiVersion: v1
kind: Secret
metadata:
  name: nsx-cert
  namespace: vmware-system-antrea
type: kubernetes.io/tls
```

```
data:
  tls.crt: LS0tLS1CRUd…LS0tLS0K
  tls.key: LS0tLS1CRUd…S0tLS0tCg==
```

What to do next

Submit the `bootstrap-config.yaml` file and the `interworking.yaml` Deployment manifest file to the Kubernetes API server to register the Antrea Kubernetes cluster to NSX.

Submit the YAML Files to the Kubernetes API Server

To register an Antrea Kubernetes cluster to NSX, you must submit the `bootstrap-config.yaml` file and the `interworking.yaml` Deployment manifest file to the Kubernetes API server.

Prerequisites

Ensure that:

- The prerequisite tasks for registering an Antrea Kubernetes cluster to NSX are completed. See Prerequisites for Registering an Antrea Kubernetes Cluster to NSX.
- Values of mandatory arguments are specified in the `bootstrap-config.yaml` file. See Edit the Bootstrap Configuration File.

Procedure

1. Run the following kubectl command to submit the `.yaml` files to the Kubernetes API server:

   ```
   $ kubectl apply -f bootstrap-config.yaml -f interworking.yaml
   ```

   Ensure that `bootstrap-config.yaml` file comes first in the command.

   This command registers the Antrea Kubernetes cluster to NSX. The register-xxx and interworking-yyy pods are deployed in the `vmware-system-antrea` namespace.

   Where: *xxx* and *yyy* are arbitrary numbers that represent instance IDs of the pods in your cluster.

2. Run the following kubectl command to view the list all pods in the `vmware-system-antrea` namespace:

   ```
   $ kubectl get pods -o wide -n vmware-system-antrea
   ```

   Verify that status of the register pod is `Completed` and the status of the interworking pod is `Running`.

The containers of the Management Plane Adapter and the Central Control Plane Adapter in the interworking pod now start running in the Antrea Kubernetes cluster. The resources in the Antrea Kubernetes cluster are synced with the NSX inventory.

Note After the Antrea Kubernetes cluster is registered to NSX, the Management Plane Adapter connects with the NSX Management Plane and performs a full synchronization of the Antrea Kubernetes cluster resources in the NSX inventory. The time required to do a full sync operation is directly proportional to the scale of the cluster. Thereafter, only a delta synchronization operation happens at regular predefined intervals. If the Management Plane Adapter fails due to any reason, the resources are not synced with the NSX inventory. Only after the adapter is up again, the resources in the Antrea Kubernetes cluster are compared with the existing objects in the NSX inventory and the difference (delta) is synchronized.

3 Perform this step only when your Antrea Kubernetes cluster uses Kubernetes version ≤ 1.20.

a Run the following kubectl command to register the Antrea Controller webhook on namespace creating events.

```
$ kubectl apply -f ns-label-webhook.yaml
```

You can find this Webhook definition file in the `antrea-interworking-version.zip` file, which you downloaded from the **Download VMware Antrea** page.

b Restart the Antrea Controller pod.

```
kubectl rollout restart deployment antrea-controller -n kube-system
```

This command deletes the existing Antrea Controller pod and creates a new Antrea Controller pod.

c Verify whether the new Antrea Controller pod is running.

```
$ kubectl get pod -l component=antrea-controller -n kube-system
```

What to do next

View the inventory of Antrea Kubernetes cluster resources, such as pods, namespaces, Antrea Network Policies, Antrea Cluster Network Policies, and other resources in the NSX Manager UI.

For more information, see:

- View Details of Namespaces in an Antrea Kubernetes Cluster
- View Details of an Antrea Kubernetes Cluster

Viewing Inventory of an Antrea Kubernetes Cluster in NSX Manager

In an NSX environment, you can view the resources of the registered Antrea Kubernetes clusters in **read-only** mode. You cannot modify or delete the Antrea Kubernetes cluster resources in NSX.

For example, in the NSX environment, you cannot modify, add, or remove tags (labels) that are attached to pods.

Note After the Antrea Kubernetes cluster is registered to NSX, the Management Plane Adapter connects with the NSX Management Plane and performs a full synchronization of the Antrea Kubernetes cluster resources in the NSX inventory. The time required to do a full sync operation is directly proportional to the scale of the cluster. Thereafter, only a delta synchronization operation happens at regular predefined intervals. If the Management Plane Adapter fails due to any reason, the resources are not synced with the NSX inventory. Only after the adapter is up again, the resources in the Antrea Kubernetes cluster are compared with the existing objects in the NSX inventory and the difference (delta) is synchronized.

Important In a multi-tenant NSX environment, Kubernetes cluster resources are not exposed to the project inventory. In other words, a project administrator cannot view and manage Antrea Kubernetes cluster resources inside a project. You must switch to the **Default** view (default space) to view and manage the inventory of registered Antrea Kubernetes clusters.

View Details of an Antrea Kubernetes Cluster

You can view the list of all Antrea Kubernetes clusters that are registered to NSX, and filter this list based on several filter criteria, such as External ID, Cluster Name, CNI Type, Type of Container Cluster, and some more criteria.

When the inventory of Antrea Kubernetes clusters is available in NSX inventory, cluster administrators can do basic analysis and diagnostic tasks in the NSX Manager UI.

For example:

- View the Kubernetes services that are running in the cluster.
- Verify whether a specific Kubernetes service is up or down.
- Verify the health status of Antrea Agent on each node in the Kubernetes cluster.
- View the pods that are running in the cluster.
- Check whether any pods in the cluster are down.
- Determine which pods in the cluster are associated with a specific Kubernetes service.
- Read the specifications of Antrea cluster network policies (ACNP).

Prerequisites

Antrea Kubernetes clusters must be registered to NSX.

For a detailed information about registering an Antrea Kubernetes cluster, see Registering an Antrea Kubernetes Cluster to NSX.

Procedure

1 From your browser, log in to an NSX Manager at `https://nsx-manager-ip-address`.

2 Navigate to **Inventory > Containers > Clusters**.

A list of all container clusters in the NSX inventory is displayed in the table.

Important NSX Manager UI fetches the information about registered Antrea Kubernetes clusters when you start the NSX Manager application in the browser. If the application UI is already open, it does not fetch the Antrea Kubernetes cluster registration information automatically. This behavior is expected and per the current UI design.

If you have no Antrea Kubernetes clusters registered to NSX and you register the first cluster while the UI is open, a forced browser refresh is required after navigating to the **Clusters** page. This manual browser refresh is required only once after the first cluster is registered, and not every time after a new Antrea Kubernetes cluster is registered to NSX. If there are existing Antrea Kubernetes clusters registered to NSX, you do not have to force refresh the browser for fetching the newly added clusters. Clicking the **Refresh** link on the **Clusters** page can fetch the updated list.

3 To view the list of only Antrea Kubernetes clusters, filter the table by **CNI Type** as **Antrea**.

4 To view the information about all nodes, Kubernetes services, and pods in a specific Antrea Kubernetes cluster, click the hyperlinked number in the respective columns.

5 Expand a row in the table to view more details about a specific Antrea Kubernetes cluster.

For example, view the following details:

- Type of infrastructure for the Antrea Kubernetes cluster (example: vSphere, AWS, Azure, Google, VMC, and so on).
- Antrea version installed in the Kubernetes cluster.
- Number of Antrea cluster network policies in the Kubernetes cluster.
- Specifications (YAML manifest) of the Antrea cluster network policies. These policies are cluster-scoped.
- Number of namespaces in the Kubernetes cluster and the details of each namespace.
- NodePortLocal range. Default range is 61000–62000.
- NodePort range. Default range is 30000–32767.
- Number of Antrea Egresses, egress IPs, and specifications (YAML manifest) of the Egress resources.
- Number of Antrea IP pools and specifications (YAML manifest) of the IP pools.

This list does not mention all the inventory details that you can view for an Antrea Kubernetes cluster. Check the **Clusters** page in the NSX Manager UI to know more.

Note

- For Antrea Kubernetes clusters, the **Networking Status** column displays `Not Applicable`. Currently, this column is used only for container clusters that use NSX Container Plugin (NCP) as the CNI.
- When NSX discovers labels on Kubernetes cluster resources, such as nodes, namespaces, pods, services, and so on, these labels are always displayed with a `dis:k8s` prefix in the NSX inventory. Remember that a label in Kubernetes maps to a tag in NSX, whereas a key in Kubernetes maps to a scope in NSX.

View Details of Namespaces in an Antrea Kubernetes Cluster

You can view the list of all namespaces that are created in Antrea Kubernetes clusters, and filter this list based on several criteria, such as ID, Cluster Name, CNI Type, Type of container cluster, and some more criteria.

When the inventory of Antrea Kubernetes clusters is available in NSX inventory, cluster administrators can do basic analysis and diagnostic tasks in the NSX Manager UI.

For example:

- View the number of pods running in each namespace, the number of pods that are down, and so on.
- View the number of Antrea network policies and Kubernetes network policies.
- Read the specifications of Antrea network policies and Kubernetes network policies.
- View the ingress rules that are associated with a namespace.
- View the Kubernetes services that are running in the namespace.

Prerequisites

Antrea Kubernetes clusters must be registered to NSX.

For a detailed information about registering an Antrea Kubernetes cluster, see Registering an Antrea Kubernetes Cluster to NSX.

Procedure

1. From your browser, log in to an NSX Manager at `https://nsx-manager-ip-address`.
2. Navigate to **Inventory > Containers > Namespaces**.

   A list of all namespaces across all container clusters is displayed in the table.
3. To view the list of namespaces created only in Antrea Kubernetes clusters, filter the table by **CNI Type** as **Antrea**.

4 Expand a row in the table to view more details about the namespace.

For example, view the following details:

- Number of ingress rules and the details of each ingress rule in the namespace.
- Number of network policies in the namespace. These policies refer to Antrea network policies and Kubernetes network policies, which have their scope limited to a namespace.
- Specifications (YAML manifest) of the network policies that are associated with the namespace.
- Number of labels and the details of each label in the namespace.
- Number of Kubernetes Gateways in the namespace and the specifications (YAML manifest) of the Gateway resources.

This list does not mention all the inventory details that you can view for a namespace. Check the **Namespaces** page in the NSX Manager UI to know more.

Note

- For Antrea Kubernetes clusters, the **Networking Status** column displays `Not Applicable`. Currently, this column is used only for container clusters that use NSX Container Plugin (NCP) as the CNI.
- When NSX discovers labels on Kubernetes cluster resources, such as nodes, namespaces, pods, services, and so on, these labels are always displayed with a `dis:k8s` prefix in the NSX inventory. Remember that a label in Kubernetes maps to a tag in NSX, whereas a key in Kubernetes maps to a scope in NSX.

Monitor Health Status of an Antrea Kubernetes Cluster

The overall health status of the Antrea Kubernetes cluster is aggregated or computed from the status of the various Antrea components and displayed in the NSX Manager UI.

The following Antrea components expose their health status to NSX Manager:

- Antrea Controller
- Antrea Agent
- Management Plane Adapter
- Central Control Plane Adapter

The Monitoring CustomResourceDefinition (CRD) objects in Antrea report the statuses of these Antrea components to NSX Manager after a predefined period called a heartbeat interval. This heartbeat interval is configurable for each cluster. The default value is 60 seconds. The permitted range of values is 60 seconds through 600 seconds. You can modify the default interval by running an NSX API.

To read the heartbeat configuration of a specific Antrea Kubernetes cluster, run the following NSX GET API:

```
GET https://{nsx-mgr-ip}/policy/api/v1/infra/sites/{site-id}/enforcement-points/
{enforcementpoint-id}/cluster-control-planes/{cluster-name}/heartbeat-config
```

To update the heartbeat configuration of a specific Antrea Kubernetes cluster, run the following NSX PUT API:

```
PUT https://{nsx-mgr-ip}/policy/api/v1/infra/sites/{site-id}/enforcement-points/
{enforcementpoint-id}/cluster-control-planes/{cluster-name}/heartbeat-config
{
    "report_interval": 120,
    ""_revision": 0
}
```

The PUT API body shows 120 as the sample report interval. You can specify any integer value from 60 through 600. The unit of reporting interval is seconds.

The `_revision` parameter describes the current revision of the `heartbeat-config` resource. PUT operation must include the current revision of this resource, which you can obtain by submitting the GET API. If the revision provided in a PUT request is missing or stale, the update operation is rejected.

For a detailed information about all the parameters in the API, including examples of the GET and PUT API responses, see the *NSX API Guide*.

Note If the Antrea components do not send a heartbeat to NSX Manager, the status of that component is shown as Unknown. This status means that status monitoring is not working. However, container networking is working on the node. Existing NSX security policies are still enforced on the pods, but if any new security policies are applied, they are not enforced on the pods.

The following procedure explains the steps to view these statuses in NSX Manager:

- Overall health status of an Antrea Kubernetes cluster.
- Health status of Antrea Agent on each node of the cluster.

Prerequisites

Antrea Kubernetes clusters are registered to NSX.

Procedure

1 From your browser, log in to an NSX Manager at `https://nsx-manager-ip-address`.

2 View the overall health status of an Antrea Kubernetes cluster.

   a Navigate to **System > Fabric > Nodes > Container Clusters > Antrea**.

   A list of all registered Antrea Kubernetes clusters is displayed. The **Status** column displays the overall health status of each cluster.

   b Click Up or Down in the **Status** column to view more details in a pop-up window.

   The overall health status of the Antrea Kubernetes cluster is computed from the status of the following Antrea components:

   - Antrea Controller
   - Management Plane Adapter
   - Central Control Plane Adapter

   If the status of any one component or all three Antrea components is Down, the overall cluster status is Down. Click **Failed/Down** in the pop-up window to view the error message. The cluster status is Up only when the statuses of all the three Antrea components are Up.

   The pop-up window also displays the total number of Antrea Agents that are Healthy, Failed, and Degraded. If an Antrea Agent is degraded, it means that the container networking on the node is working. However, new security policies might not be enforced correctly on the node. If an Antrea Agent has failed, it means that the container networking on that node is not working.

   To view the status of each individual node in the Antrea Kubernetes cluster, check the Antrea Agent status on each node, as explained in the next step.

3 Check the health status of Antrea Agent on each node of the Antrea Kubernetes cluster.

   a Navigate to **Inventory > Containers > Clusters**.

   A list of all container clusters in the NSX inventory is displayed.

   b Filter this list by **CNI Type** as **Antrea**.

   c Click the hyperlinked number in the **Nodes** column.

   The **Nodes** window opens. The **Agent Status** column shows whether the Antrea Agent on the node is Up or Down. The **Agent Status** column does not show Degraded as one of the statuses.

Trace the Path of a Packet with Antrea Traceflow

You can start a Traceflow session in NSX Manager to trace the path of packet in an Antrea Kubernetes cluster. Antrea Traceflow currently supports tracing the path of only Unicast traffic. Broadcast and Multicast traffic are not supported.

Note In a multi-tenant NSX environment, the Antrea Traceflow feature is currently not supported under projects. You must use this feature in the **Default** view (default space) of the NSX environment.

The source of a Traceflow session must be a pod, whereas the destination can be a pod or a service in the same Antrea Kubernetes cluster. You can trace the path of a packet for the following types of traffic in an Antrea Kubernetes cluster:

- Pod to pod traffic on the same node (intra-node traffic)
- Pod to pod traffic between nodes (inter-node traffic)
- Pod to service traffic on the same node
- Pod to service traffic between nodes
- Pod to an arbitrary IP address

Note Currently, Antrea Traceflow does not support tracing the path of a packet that travels between Antrea Kubernetes cluster and NSX overlay network. However, some workaround is possible. For more information, see Workaround to Trace the Packet Flow Between an Antrea Kubernetes Cluster and an NSX Overlay Network.

Traceflow injects a test packet into the Antrea Kubernetes cluster network and monitors the flow of the packet. As the packet flows from the source to destination, observations are collected from various components along the path of the packet. These observations are displayed in the Traceflow output, which shows the various components in the path of the packet.

Prerequisites

Antrea Kubernetes cluster is registered to NSX.

Procedure

1 From your browser, log in to an NSX Manager at `https://nsx-manager-ip-address`.

2 Navigate to **Plan & Troubleshoot > Traffic Analysis**.

Important NSX Manager UI fetches the information about registered Antrea Kubernetes clusters when you start the NSX Manager application in the browser. If the application UI is already open, it does not fetch the Antrea Kubernetes cluster registration information automatically. This behavior is expected and per the current UI design. If you have registered the first Antrea Kubernetes cluster after the NSX Manager application is opened, ensure that you refresh the browser after navigating to the **Traceflow Analysis** page. A manual refresh ensures that the **Antrea Traceflow** tab is visible in the UI when you reach step 4 of this procedure.

This manual browser refresh is required only once, and not every time after a new Antrea Kubernetes cluster is registered to NSX.

3 In the **Traceflow** card, click **Get Started**.

4 Click the **Antrea Traceflow** tab.

This tab is available only when at least one Antrea Kubernetes cluster is registered to NSX. If this tab is not visible, refresh the browser.

5 Specify the configuration settings of the Traceflow session.

Field	Description
Cluster	Select an Antrea Kubernetes cluster from the drop-down menu. Alternatively, enter the first few characters of the cluster name to filter the list, and then select the cluster.
IP Address	Select either **IPv4** or **IPv6**.
Protocol Type	Select any one protocol type: ICMP, TCP, UDP. For ICMP: ■ Optional: Enter the ICMP identifier. Default is 0. ■ Optional: Enter the ICMP sequence number. Default is 0. ■ Optional: Enter the TTL. Default is 64. For TCP, UDP: ■ Optional: Enter the source port number. Default is 0. ■ Optional: Enter the destination port number. Default is 0. ■ (Only for TCP): Add TCP flags, if required. SYN flag is set by default. This flag is required for the Antrea Traceflow. ■ Optional: Enter the TTL. Default is 64.

Field	Description
Source	Only **Pod** is supported as the source of an Antrea Traceflow. **Note** Pods that use the host network are currently not supported as the source of an Antrea Traceflow. This limitation is a known behavior. For example, Antrea Agent pods are not supported in the source of an Antrea Traceflow. If you know the pod name, select it directly from the **Pod** drop-down menu. Otherwise, filter the list of pods by doing these steps: a Enter the first few characters of the node name, or select a value from the **Node** drop-down menu. b Enter the first few characters of the namespace, or select a value from the **Namespace** drop-down menu.
Destination	a Select the type of destination: **Pod**, **Service**, **IP Address**. b If you select **Pod** or **Service** as the destination, use the **Node** and **Namespace** drop-down menus to filter the list of pods or services. c If you select **IP Address** as the destination, enter an IP address.

6 Click **Trace**.

Results

The Traceflow observations are displayed in a tabular format. For each observation, the table shows the following information.

Observation Type

This column takes the following values.

Observation Type	Description
Delivered	The packet is delivered to destination Pod or Service properly.
Dropped	The packed is dropped by a network policy.
Received	The packet is received from another node in the Kubernetes cluster.
Forwarded	The packet is forwarded to the next logical node or a Kubernetes cluster object.

Component

This column shows the components that the test packet had encountered on its path from the source to the destination. Sample component values are: IngressRule, EgressRule, SpoofGuard, Classification, Output, and so on.

Click the component name to view more information in a pop-up window.

Timestamp

The date and time for each observation.

Workaround to Trace the Packet Flow Between an Antrea Kubernetes Cluster and an NSX Overlay Network

In NSX 4.1, Antrea Traceflow does not support tracing the path of a packet that travels between Antrea Kubernetes cluster and NSX overlay network. However, some workaround is possible, as explained in this documentation.

Before following the workflow instructions in this documentation, read this important note carefully to ensure that the configuration of the **Protocol Type** and **Destination Port** fields are done correctly in the NSX Traceflow and Antrea Traceflow sessions.

Important By default, the **Protocol Type** of a Traceflow session is set to **`ICMP`**. When you change it **`TCP`**, the **Destination Port** is set to **`0`** by default. Remember to specify the port number of the Kubernetes service or the NSX Load Balancer virtual server. This is because the K8s service and NSX Load Balancer virtual server listen on a specific TCP port. In short, you must specify all the information (IP address, Protocol Type, and Destination Port) accurately in the Traceflow session to ensure that the injected packet is processed by the K8s service or the NSX Load Balancer.

Trace the Path of a Packet from NSX VM to K8s Pod

If you want to trace the path of a packet from a VM in an NSX overlay network to a pod workload in an Antrea Kubernetes cluster, do the following steps:

1. Navigate to **Plan & Troubleshoot > Traffic Analysis**, and in the **Traceflow** card, click **Get Started**.
2. In the **Source** area, select **Virtual Machine** from the **Type** drop-down menu.
3. In the **Destination** area, select **IP-Mac**. Click the **Layer 3** option, and then enter any one of these IP addresses:
   - Kubernetes Ingress IP address
   - LoadBalancer Service IP address
   - Node IP address
   - Pod IP address (if pod IP is routable on the underlying network)

NSX Traceflow will trace the path of the packet till it leaves the NSX overlay network.

Currently, the following components are not suported in the **Source** of an Antrea Traceflow:

- Kubernetes Service (LoadBalancer, NodePort)
- Kubernetes Gateway
- Kubernetes Ingress

Therefore, you cannot trace the path of the packet after it reaches the Antrea Kubernetes cluster.

Trace the Path of a Packet from K8s Pod to NSX VM

If you want to trace the path of a packet from a pod in an Antrea Kubernetes cluster to a VM in an NSX overlay network, do the following steps:

1 Navigate to **Plan & Troubleshoot > Traffic Analysis**, and in the **Traceflow** card, click **Get Started**.

2 Click the **Antrea Traceflow** tab.

3 In the **Source** area, select **Pod** from the **Type** drop-down menu.

4 In the **Destination** area, select **IP address** from the **Type** drop-down menu, and then enter the IP address of the NSX Load Balancer VM.

Antrea Traceflow will trace the path of the packet till it leaves the Antrea Kubernetes cluster.

As explained in the important note earlier in this documentation, esnure that you all the information (IP address, Protocol Type, and Destination Port) accurately in the Antrea Traceflow session

5 Now, you use the NSX Traceflow feature to trace the packet from a tier-0 gateway interface to the destination VM in the NSX overlay network. Configure the **Source** in the NSX Traceflow as follows:

- **Type**: Port/Interface
- **Attachment**: Edge Uplink
- **Port**: Select the tier-0 interface from the drop-down menu.

Configure the **Destination** in the NSX Traceflow as follows:

- **Type**: IP-Mac
- **Layer**: Layer-3 (Network Layer)
- **IP Address**: Enter the IP address of the NSX Load Balancer VM.

The trace result will display how the NSX Load Balancer chooses a VM as the destination, and what path the packet has taken from the NSX Load Balancer to the destination VM.

Antrea Groups

You can create Antrea groups and use them as sources or destinations in distributed firewall policies to protect traffic between pods within an Antrea Kubernetes cluster.

Note In a multi-tenant NSX environment, Antrea groups are currently not supported under projects. You can add Antrea groups only in the **Default** view (default space) of the NSX environment.

Antrea groups feature is supported only when your NSX environment has one or more Antrea Kubernetes clusters registered to it. When registered Antrea Kubernetes clusters are detected, NSX Manager shows a separate group type called **Antrea** on the **Add Group** page of the UI. You must select this group type to add Antrea groups.

If your goal is to protect traffic between Antrea Kubernetes clusters and VMs in the NSX overlay network, you must create **Generic** groups with Kubernetes member types in dynamic membership criteria and use these groups in firewall rules. For more information, see Generic Groups with Kubernetes Member Types in Dynamic Membership Criteria.

An Antrea group can include static IP addresses, membership criteria, or both. IP addresses can be pod IP or service IP addresses.

When an Antrea group contains membership criteria, the effective members computed by that membership criteria can only be pods.

Note

- Effective members are computed for Antrea groups only when the Antrea groups are used in distributed firewall rules.

  When you add Antrea groups with membership criteria, but do not use these groups in any of the distributed firewall rules, the effective members of these Antrea groups are not computed or evaluated in NSX. In other words, the **Effective Members** page of these Antrea groups is empty.
- When you add static IP addresses in Antrea groups, effective members are currently not displayed in the UI, regardless of whether the groups are used in distributed firewall rules.

To add membership criteria in an Antrea group, the following member types are currently supported:

- Namespace
- Service
- Pod

Overview of Membership Criteria

You can add Antrea groups with a single membership criterion or multiple criteria. A membership criterion consists of one or multiple conditions. A condition in a membership criterion consists of the following properties:

- Member type
- Either name of the member type or a tag that is attached to the member type
- Tag operator and value (only when tag is used)
- Scope operator and value (only when tag is used)

The conditions in a membership criterion can use the same member type or a mix of different member types. For example, if a membership criterion consists of three conditions, the first two conditions can be based on Pod, whereas the third condition can be based on Namespace. However, some restrictions exist for adding multiple conditions in a membership criterion. See the "Supported and Unsupported Features" section later in this topic.

By default, NSX uses the logical AND operator after each condition in a membership criterion. Other logical operators are not supported to join conditions in a membership criterion.

Examples:

Membership Criteria	Description
Criterion 1: Pod Tag Equals App Scope Equals Servers	Membership criterion consists of only a single condition that is based on Pod. Multiple conditions are not used. The effective members of this Antrea group include all pods with the App tag.
Criterion 2: Pod Tag Equals App Scope Equals Servers Pod Tag Equals DB Scope Equals Servers Pod Tag Equals Web Scope Equals Servers	Membership criterion consists of three conditions. All conditions in the criterion are based on Pod. The effective members of this Antrea group include all pods with the App, DB, and Web tags.
Criterion 3: Namespace Name Equals Production Service Name Equals Cache	Membership criterion consists of two conditions with a mix of Namespace and Service. The effective members of this Antrea group include all pods that are associated with the Cache Service in the Production Namespace.

Joining Membership Criteria with OR, AND Operators

An Antrea group supports multiple membership criteria. To join the criteria, OR and AND operators are available. By default, NSX selects the OR operator to join multiple criteria. AND operator is supported between two criteria only when:

- Both criteria use the same member type.
- Both criteria use a single condition.

Examples:

- Criterion 1, Criterion 2, and Criterion 3 are all based on Pod, and they do not contain multiple conditions. In this case, Criterion 1 and Criterion 2 can be joined with either OR or AND operator. Similarly, Criterion 2 and Criterion 3 can also be joined with either OR or AND operator.
- Criterion 1 is based on Pod, whereas Criterion 2 uses two conditions: one based on Service and the other based on Namespace. In this case, only OR operator is supported for joining Criterion 1 and 2. AND operator is not allowed.

- Criterion 1 and Criterion 2 are based on Pod, whereas Criterion 3 uses two conditions: one based on Service and the other based on Namespace. In this case, Criterion 1 and Criterion 2 can be joined with either AND or OR operator. However, Criterion 2 and Criterion 3 can be joined only with OR operator. AND operator is not allowed.

Supported and Unsupported Features

The following table lists the member types, tag operator, and scope operator that are supported for adding membership criteria in Antrea groups.

Member Type	Object Attribute	Tag Operator	Scope Operator	Example Criteria
Namespace	Name	Equals	Not applicable	Namespace Name Equals Production
Namespace	Tag	Equals Not Equals	Equals	Namespace Tag Equals DB Scope Equals Servers
Service	Name	Not supported	Not supported	Service Name Equals Cache
Pod	Tag	Equals Not Equals	Equals	Pod Tag Equals App Scope Equals Servers

- The following tag operators are currently not supported for Namespace and Pod member types:
  - Contains
  - Starts With
  - Ends With
- In a membership criterion, a condition based on Service must be combined with a condition based on the Name attribute of a Namespace. In other words, a criterion with only the Service member type is not allowed.

  Example:

Supported	Not Supported
Criterion: Service Name Equals My-Service Namespace Name Equals Staging	Criterion: Service Name Equals My-Service

- In a membership criterion, a condition based on Service cannot be combined with a condition based on Pod. However, you can add conditions based on Service and Pod in two separate membership criteria and join them with an OR operator.

  Example:

Supported	Not Supported
Criterion 1: Service Name Equals My-Service OR Criterion 2: Pod Tag Equals DB Scope Equals Servers	Criterion: Service Name Equals My-Service Pod Tag Equals DB Scope Equals Servers

- For adding members statically in an Antrea group, only IP addresses are supported. Namespace, Pod, and Service cannot be added statically as members in an Antrea group.
- When you are adding an Antrea group, NSX shows an information message if you try to change the group type from **Antrea** to **Generic**, or from **Antrea** to **IP Addresses Only**. On confirming the change, all the membership criteria in the group are lost. Only the IP addresses are retained in the group.

  After an Antrea group is realized (saved) in NSX, you cannot change the group type. The **Generic** and **IP Addresses Only** group types are dimmed.

Workaround for Kubernetes Version ≤ 1.20

Antrea group criterion "Namespace Name Equals *Value*" works with Kubernetes version ≥ 1.21.

Kubernetes 1.21 or later automatically adds a special label to all namespaces, and criterion "Namespace Name Equals *Value*" internally uses this special label. However, for Kubernetes version ≤ 1.20, a workaround is required. You must register the Antrea Controller webhook on namespace creating events. When the Antrea Controller webhook is called, Antrea Controller adds a special label to the new namespace, so criterion "Namespace Name Equals *Value*" can use this label. For details about registering the Antrea Controller webhook, see step 3 in Submit the YAML Files to the Kubernetes API Server.

Note Antrea Controller webhook is effective only for new namespaces that you create after registering the webhook. In other words, existing namespaces, such as `kube-system` and `default` do not get the special label, and criterion "Namespace Name Equals *Value*" does not work with these namespaces.

Add an Antrea Group

You can add static IP addresses, membership criteria, or both in Antrea groups, and then use these groups as the source or destination of the distributed firewall policies that you create to protect traffic within an Antrea Kubernetes cluster.

Prerequisites

At least one Antrea Kubernetes cluster is registered to NSX.

Procedure

1 From your browser, log in to an NSX Manager at `https://nsx-manager-ip-address`.

2 Navigate to **Inventory > Groups**.

Note NSX Manager UI fetches the information about registered Antrea Kubernetes clusters when you start the NSX Manager application in the browser. If the application UI is already open, it does not fetch the Antrea Kubernetes cluster registration information automatically. This behavior is expected and per the current UI design. If you have registered the first Antrea Kubernetes cluster after the NSX Manager application is opened, ensure that you refresh the browser after navigating to the **Groups** page. A manual refresh ensures that the **Antrea** group type option is visible in the UI when you reach step 5 of this procedure.

This manual browser refresh is required only once, and not every time after a new Antrea Kubernetes cluster is registered to NSX.

3 Click **Add Group**.

4 Enter a name and optionally a description for the group.

5 Click **Set** and select **Antrea** as the group type.

An Antrea group can include membership criteria, static IP addresses, or both. Depending on your requirements, perform steps 6 or 7 or both.

6 To add a membership criterion, click **Add Criterion**.

a In the **Criterion** pane, select the member type on which you want to define the condition.

The supported member type are: Namespace, Service, and Pod.

b Specify the properties of the condition, such as Name or Tag, Tag operator, Scope operator, as required.

c (Optional) To add more than one condition in a membership criterion, click the plus icon in the upper-right corner of the **Criterion** pane, and define the properties of the condition.

In a membership criterion, NSX joins all the conditions with the AND operator, by default. OR operator is not supported.

d (Optional) To add multiple criteria, click **Add Criterion** again.

To join membership criteria, AND and OR operators are available. By default, NSX selects the OR operator to join two criteria. AND operator is supported between two criteria only when:

- Both criteria use the same member type.
- Both criteria use a single condition.

For more information about what is supported and not supported for adding membership criteria, see Antrea Groups.

7 To add static IP addresses in the group, click **IP Addresses**, and enter IP values in the text box.

If you want to import IP values from a TXT or a CSV file, click **Actions > Import** . The values in the file must be separated with commas. The allowed values are IP addresses, IP ranges, or IP addresses in a CIDR format. You can also do a combination of both actions. That is, enter values in the text box and import values from a file. However, the total number of IP values in the text box must not exceed the maximum limit that is displayed on the **IP Addresses** tab.

8 Click **Apply**, and then click **Save**.

Results

The Antrea group is saved in NSX and the status changes to Success.

Note

- Effective members are computed for Antrea groups only when the Antrea groups are used in distributed firewall rules.

  When you add Antrea groups with membership criteria, but do not use these groups in any of the distributed firewall rules, the effective members of these Antrea groups are not computed or evaluated in NSX. In other words, the **Effective Members** page of these Antrea groups is empty.

- When you add static IP addresses in Antrea groups, effective members are currently not displayed in the UI, regardless of whether the groups are used in distributed firewall rules.

Example: Add an Antrea Group Based on Pods

Assume that you want to add an Antrea group that contains all pods running the Revenue, Sales, and Metrics financial applications across all the namespaces in the Antrea Kubernetes cluster.

Consider that the following tags are attached to pods in the Antrea Kubernetes cluster.

Tag	Scope
RevenueApp	Finance
SalesApp	Finance
MetricsApp	Finance

Create a membership criterion with three conditions based on the Pod member type as follows:

Criterion:

Pod Tag Equals RevenueApp Scope Equals Finance

Pod Tag Equals SalesApp Scope Equals Finance

Pod Tag Equals MetricsApp Scope Equals Finance

By default, NSX uses the AND operator after each condition. When this Antrea group is used in a distributed firewall rule, the effective pod members for this group are computed.

After the distributed firewall policy is realized, go to the **Add Group** page. Click **View Members** for this Antrea group, and verify that the effective pod members are displayed on the **Effective Members** page.

Generic Groups with Kubernetes Member Types in Dynamic Membership Criteria

You can create groups with Kubernetes member types (resources) in dynamic membership criteria to match traffic entering into or leaving from Antrea Kubernetes clusters.

You can then use these generic groups in distributed firewall rules or gateway firewall rules to secure traffic between VMs in the NSX environment and pods in Antrea Kubernetes clusters.

This documentation uses the phrase "Kubernetes member types" to refer to "Kubernetes resources" that you can use to configure dynamic membership criteria.

Currently, generic groups with Kubernetes member types support only dynamic membership criteria. You cannot add Kubernetes member types statically in a generic group definition.

Note To evaluate the effective members of a generic group that contains dynamic membership criteria defined using Kubernetes member types, NSX considers the resource inventory that is reported by Kubernetes clusters with Antrea CNI. The resource inventory that is reported by Kubernetes clusters with NCP CNI is ignored for evaluating the effective group members.

Kubernetes Member Types in Membership Criteria

In a generic group, Kubernetes member types are available on the **Membership Criteria** page only when at least one Antrea Kubernetes cluster is registered to your NSX environment.

Note In a multi-tenant NSX environment, Kubernetes cluster resources are not exposed to the project inventory. Therefore, inside a project, you cannot create generic groups with Kubernetes member types in dynamic membership criteria.

The following table lists the Kubernetes member types that are available in the **Default** view of NSX 4.1 onwards to add dynamic membership criteria in a generic group. Each Kubernetes member type belongs to either cluster scope or namespace scope.

Member Type	Scope
Kubernetes Cluster	Cluster
Kubernetes Namespace	Namespace
Kubernetes Node	Cluster
Kubernetes Service	Namespace
Kubernetes Ingress	Namespace
Kubernetes Gateway	Namespace

Member Type	Scope
Antrea Egress	Cluster
Antrea IP Pool	Cluster

Naming Conventions Used for Conditions With Kubernetes Member Types

The following table explains the naming conventions that are used in this documentation to represent the different conditions that you can add in dynamic membership criteria, which are based on Kubernetes member types.

Naming Convention for Condition	Meaning
Kubernetes Cluster condition	Conditions in the dynamic membership criteria are based on the Kubernetes Cluster member type.
Kubernetes Namespace condition	Conditions in the dynamic membership criteria are based on the Kubernetes Namespace member type.
Resource condition	Conditions in the dynamic membership criterion are based on any of these Kubernetes member types: ■ Kubernetes Service ■ Kubernetes Ingress ■ Kubernetes Gateway ■ Antrea Egress ■ Antrea IP Pool ■ Kubernetes Node
Cluster-scoped resource condition	Conditions in the dynamic membership criteria are based on any of these Kubernetes member types: ■ Antrea Egress ■ Antrea IP Pool ■ Kubernetes Node
Namespace-scoped resource condition	Conditions in the dynamic membership criteria are based on any of these Kubernetes member types: ■ Kubernetes Service ■ Kubernetes Ingress ■ Kubernetes Gateway

Overview of Membership Criteria with Kubernetes Member Types

A criterion can have one or more conditions. The conditions can use the same Kubernetes member type or a mix of different Kubernetes member types. However, some restrictions apply to adding multiple conditions with mixed member types in a membership criterion. See the *Restrictions for Using Kubernetes Member Types in a Criterion* section later in this documentation.

By default, NSX uses the logical AND operator after each condition in a membership criterion. Other logical operators are not supported to join the conditions in a membership criterion.

To join criteria, OR and AND operators are available. By default, NSX selects the OR operator to join two criteria. AND operator is supported between two criteria only when:

- Both criteria use the same Kubernetes member type.
- Both criteria use a single condition.

The following restrictions apply to adding multiple conditions:

- A maximum of five conditions with the same Kubernetes member type is supported in a single membership criterion. For example, in a criterion, you can add a maximum of five Kubernetes Service conditions.
- A maximum of 15 conditions with mixed Kubernetes member types are supported in a single membership criterion. For example, in a criterion, you can add a maximum of 15 conditions with a mix of Kubernetes Namespace conditions and Kubernetes Ingress conditions.
- A maximum of 35 conditions with mixed member types are supported in a generic group.

A group can have a maximum of five membership criteria. However, the total number of criteria that you can add in a generic group is determined by the number of conditions in each criterion. See the following examples.

Example 1

A generic group with three membership criteria and a total of 35 conditions:

- Criterion 1 has 15 conditions with mixed member types.
- Criterion 2 has 15 conditions with mixed member types.
- Criterion 3 has 5 conditions with the same member type.

Example 2

A generic group with four membership criteria and a total of 35 conditions:

- Criterion 1 has 15 conditions with mixed member types.
- Criterion 2 has 14 conditions with mixed member types.
- Criterion 3 has four conditions with the same member type.
- Criterion 4 has two conditions with the same member type.

Example 3

A generic group with five membership criteria and a total of 22 conditions:

- Criterion 1 has 10 conditions with mixed member types.
- Criterion 2 has three conditions with the same member type.
- Criterion 3 has four conditions with the same member type.
- Criterion 4 has three conditions with the same member type.
- Criterion 5 has two conditions with mixed member types.

Because this group has reached the limit of five criteria, you cannot add another membership criterion. However, you can add more conditions, if required, in any of the five criteria until you don't exceed the following upper limits mentioned earlier:

- A maximum of five conditions with the same member type in a single criterion.
- A maximum of 15 conditions with mixed member types in a single criterion.
- A total of 35 conditions in the generic group.

Restrictions for Using Kubernetes Member Types in a Criterion

The following table provides a high-level summary of the restrictions or validations that apply to using Kubernetes member types in a single membership criterion. For examples of validations, see the *Validations for Dynamic Grouping with Kubernetes Member Types* section later in this documentation.

Member Type	Restrictions for Using the Member Type in a Criterion	Supported Attributes	Tag Operator	Scope Operator
Kubernetes Cluster	Cannot be used alone in a criterion. Only one cluster condition is allowed in a criterion. Must be mixed with at least one Kubernetes resource condition. Optionally, can be mixed with Kubernetes Namespace conditions and Kubernetes resource conditions.	Name	Not Supported	Not Supported
Kubernetes Namespace	Cannot be used alone in a criterion. Cannot be mixed with cluster-scoped resource conditions. Must be mixed with namespace-scoped resource conditions. Optionally, can be mixed with a Kubernetes Cluster condition.	Name Tag	Equals - one tag can be selected	Equals
Antrea Egress	Can be used alone in a criterion. Optionally, can be mixed with a Kubernetes Cluster condition.	Name Tag	Equals - one tag can be selected	Equals
Antrea IP Pool	Can be used alone in a criterion. Optionally, can be mixed with a Kubernetes Cluster condition.	Name Tag	Equals - one tag can be selected	Equals
Kubernetes Ingress	Can be used alone in a criterion. Optionally, can be mixed with only Kubernetes Namespace conditions or only Kubernetes Cluster condition, or both.	Name Tag	Equals - one tag can be selected	Equals

Member Type	Restrictions for Using the Member Type in a Criterion	Supported Attributes	Tag Operator	Scope Operator
Kubernetes Gateway	Can be used alone in a criterion. Optionally, can be mixed with only Kubernetes Namespace conditions or only Kubernetes Cluster condition, or both.	Name Tag	Equals - one tag can be selected	Equals
Kubernetes Service	Can be used alone in a criterion. Optionally, can be mixed with only Kubernetes Namespace conditions or only Kubernetes Cluster condition, or both.	Name Tag	Equals - one tag can be selected	Equals
Kubernetes Node	Can be used alone in a criterion. Only a single node condition is allowed. Optionally, can be mixed with a Kubernetes Cluster condition. Cannot be mixed with Kubernetes Namespace conditions or Kubernetes resource conditions.	IP Address Pod CIDR	Not supported See Note after this table.	Not supported

Note When you create a generic group with Kubernetes Node member type by using the API, only the Matches operator is allowed. This operator can take only the * value. The * wildcard character will match all the nodes in the K8s cluster (if the Kubernetes Node condition is mixed with a Kubernetes Cluster condition), or it will match nodes across all clusters (if the Kubernetes Node condition is used alone).

Validations for Dynamic Grouping with Kubernetes Member Types

Validation 1

A membership criterion can have a maximum of one Kubernetes Cluster condition. To match a single Kubernetes cluster by name, use the Equals operator and enter the cluster name.

Note Kubernetes cluster name must be unique.

If you want to match multiple clusters in a single Kubernetes Cluster condition, you can use one of these operators:

- In
- Starts With
- Ends With

Example:

Matches Single K8s Cluster	Matches Multiple K8s Clusters
Criterion: Kubernetes Cluster Name Equals ClusterA	Criterion: Kubernetes Cluster Name In ClusterA,ClusterB,ClusterC A maximum of five comma-separated values are allowed. The values must not be separated with spaces.

Validation 2

A membership criterion with a Kubernetes Cluster condition can be mixed with any one of the Kubernetes resource conditions. If you add Kubernetes Namespace conditions too in the same criterion, then resource conditions must be limited only to namespace-scoped resource conditions.

Examples:

- A membership criterion with a Kubernetes Cluster condition and three Kubernetes Ingress conditions is valid.
- A membership criterion with a Kubernetes Cluster condition and two Antrea Egress conditions is valid.
- A membership criterion with a Kubernetes Cluster condition and three Antrea IP Pool conditions is valid.
- A membership criterion with a Kubernetes Cluster condition, one Kubernetes Namespace condition, and three Kubernetes Gateway conditions is valid.

Validation 3

A membership criterion must include at least one Kubernetes resource condition. A membership criterion is invalid if it contains:

- Only Kubernetes Cluster condition
- Only Kubernetes Namespace condition
- Only Kubernetes Cluster condition and Kubernetes Namespace condition

Examples:

- A membership criterion with one Kubernetes Cluster condition, one Kubernetes Namespace condition, and one Kubernetes Ingress condition is valid.
- A membership criterion with one Kubernetes Cluster condition, two Kubernetes Namespace conditions, and three Kubernetes Ingress conditions is valid.

Validation 4

A membership criterion can have only a single Kubernetes Node condition. Optionally, a Kubernetes Node condition can be mixed with only a Kubernetes Cluster condition.

A Kubernetes Node condition cannot be mixed with Kubernetes Namespace conditions or Kubernetes resource conditions.

A membership criterion with only a Kubernetes Node condition can stand alone. However, a group with only a Kubernetes Node condition will match nodes of all Antrea Kubernetes clusters that are registered to NSX.

Tag operator and Scope operator are currently not supported for the Kubernetes Node condition.

Kubernetes Node condition supports the following two properties.

Property	Description
IP Address	Matches the internal IP address of all nodes of specified Antrea Kubernetes clusters.
Pod CIDR	Matches the Pod CIDR of all nodes of specified Antrea Kubernetes clusters.

Validation 5

If a membership criterion includes Kubernetes Cluster condition and Kubernetes Namespace conditions, then it must include at least one namespace-scoped Kubernetes resource condition. You cannot mix any of the following cluster-scoped Kubernetes resource conditions in the same criterion:

- Antrea Egress
- Antrea IP Pool
- Kubernetes Node

Examples:

- A membership criterion with one Kubernetes Cluster condition, two Kubernetes Namespace conditions, and Kubernetes Gateway conditions is valid.
- A membership criterion with one Kubernetes Cluster condition, four Kubernetes Namespace conditions, and three Kubernetes Service conditions is valid.
- A membership criterion with one Kubernetes Cluster condition, one Kubernetes Namespace condition, and one Antrea Egress condition is invalid because Antrea Egress is a cluster-scoped resource.

Validation 6

A membership criterion must include at least one Kubernetes resource condition. A resource condition can stand alone in a criterion. However, if you add multiple resource conditions in a criterion, then all resource conditions must be of the same member type.

Note Kubernetes Cluster condition and Kubernetes Namespace conditions are used for defining the scope of a criterion. They are not Kubernetes resource conditions, and hence they are not limited by this validation rule.

Examples:

- A membership criterion with five Kubernetes Service conditions is valid.

- A membership criterion with one Kubernetes Cluster condition, three Kubernetes Namespace conditions, and four Kubernetes Service conditions is valid.
- A membership criterion with one Kubernetes Cluster condition, three Kubernetes Namespace conditions, four Kubernetes Service conditions, and three Kubernetes Ingress conditions is invalid. The reason is that you mixed resource conditions of two different member types (Kubernetes Service and Kubernetes Ingress) in the same criterion.

  However, you can create separate criteria with a resource condition based on a single member type and then join both criteria with an OR operator, as shown below:

  Criterion 1:

  One Kubernetes Cluster condition + three Kubernetes Namespace conditions + four Kubernetes Service conditions

  OR

  Criterion 2:

  One Kubernetes Cluster condition + three Kubernetes Namespace conditions + three Kubernetes Ingress conditions

Validation 7

In a single membership criterion, conditions based on NSX member types cannot be mixed with conditions based on Kubernetes member types. However, you can have a group with one criterion based on only NSX members types, and other criterion based on only Kubernetes member types, and then join both criteria with an OR operator.

Example:

Valid	Invalid
Criterion 1: Virtual Machine conditions OR Criterion 2: Kubernetes Cluster condition + Kubernetes Gateway conditions	Criterion: NSX Segment condition + Segment Port condition AND Kubernetes Cluster condition + Kubernetes Gateway conditions

Effective Members for Conditions with Kubernetes Member Types

From an NSX point of view, the effective members for groups with Kubernetes member types are either discrete IP addresses, IP ranges, or list of IP addresses and ports.

The following table provides some examples.

Example Group Definition	Use Case	Description	Effective Members
Kubernetes Cluster condition AND Kubernetes Node condition (based on IP address)	Traffic from Antrea Kubernetes cluster to NSX	Matches all node IP addresses in specified clusters	IP addresses
Kubernetes Cluster condition AND Kubernetes Node condition (based on Pod CIDR)	Traffic from Antrea Kubernetes cluster to NSX	Matches Pod CIDRs of all nodes in specified clusters	Pod CIDRs
Kubernetes Cluster condition AND Antrea Egress condition (based on name)	Traffic from Antrea Kubernetes cluster to NSX	Matches Egress by name in specified clusters	Egress IP addresses
Kubernetes Cluster condition AND Antrea Egress condition (based on tag) AND More Antrea Egress conditions (based on tag)	Traffic from Antrea Kubernetes cluster to NSX	Matches Egress by tags in specified clusters.	Egress IP addresses
Kubernetes Cluster condition AND Antrea IP Pool condition (based on name)	Traffic from Antrea Kubernetes cluster to NSX Traffic from NSX to Antrea Kubernetes cluster	Matches IP pool by name in specified clusters	IP ranges
Kubernetes Cluster condition AND Antrea IP Pool condition (based on tag) AND More Antrea IP Pool conditions (based on tags)	Traffic from Antrea Kubernetes cluster to NSX Traffic from NSX to Antrea Kubernetes cluster	Matches IP pools by tags in specified clusters	IP ranges
Kubernetes Cluster condition AND Kubernetes Namespace condition AND Kubernetes Ingress condition (based on name)	Traffic from NSX to Antrea Kubernetes cluster	Matches Ingress by name in specified clusters and namespace	Ingress IP addresses
Kubernetes Cluster condition AND Kubernetes Namespace condition AND Kubernetes Gateway condition (based on name)	Traffic from NSX to Antrea Kubernetes cluster	Matches Gateway by name in specified clusters and namespace	Gateway IP addresses
Kubernetes Cluster condition AND Kubernetes Namespace condition AND Kubernetes Ingress condition (based on tag) AND More Kubernetes Ingress conditions (based on tags)	Traffic from NSX to Antrea Kubernetes cluster	Matches Ingress by tags in specified clusters and namespace	Ingress IP addresses

Example Group Definition	Use Case	Description	Effective Members
Kubernetes Cluster condition AND Kubernetes Namespace condition AND Kubernetes Gateway condition (based on tag) AND More Kubernetes Gateway conditions (based on tags)	Traffic from NSX to Antrea Kubernetes cluster	Matches Gateway by tags in specified clusters and namespace	Gateway IP addresses
Kubernetes Cluster condition AND Kubernetes Namespace condition AND Kubernetes Service condition (based on tag and type=LoadBalancer) AND More Kubernetes Service conditions (based on tags and type=LoadBalancer)	Traffic from NSX to Antrea Kubernetes cluster	Matches Service (LoadBalancer) by tags in specified clusters and namespace	LoadBalancer Ingress IP addresses
Kubernetes Cluster condition AND Kubernetes Namespace condition AND Kubernetes Service condition (based on name,type=ClusterIP and NodePortLocal feature enabled)	Traffic from NSX to Antrea Kubernetes cluster	Matches Service (ClusterIP with NodePortLocal enabled) by name in specified clusters and namespace.	Node IP addresses, NodePortLocal range

Distributed Firewall Policies for Securing Traffic Within an Antrea Kubernetes Cluster

You can create Antrea groups in your NSX environment and use these groups in the distributed firewall policies (security policies) to secure traffic between pods within an Antrea Kubernetes cluster.

Note In a multi-tenant NSX environment, Antrea groups are currently not supported under projects. Therefore, you cannot create security policies under projects to secure traffic between pods within an Antrea Kubernetes cluster. You must create the security policies in the **Default** view (default space) of the NSX environment.

An NSX security policy can be applied to multiple Antrea Kubernetes clusters. However, the distributed firewall (DFW) policy can secure traffic between pods within a single Antrea Kubernetes cluster. Pod-to-pod traffic between Antrea Kubernetes clusters is currently not protected.

When an NSX security policy is applied to one or more Antrea Kubernetes clusters, the Antrea network plug-in enforces this security policy at the Antrea Controller of each Kubernetes cluster. In other words, the enforcement point of the security policy is the Antrea Controller of each Antrea Kubernetes cluster.

If your goal is to protect traffic between pods in an Antrea Kubernetes cluster and VMs on hosts in the NSX environment, see Firewall Policies for Securing Traffic Between Antrea Kubernetes Clusters and VMs in an NSX Network.

Security Policy Features Supported for Antrea Kubernetes Clusters

- Only Layer 3 and 4 security policies can be applied to Antrea Kubernetes clusters. Rules in the following firewall categories are supported: Emergency, Infrastructure, Environment, and Application.
- Sources, Destinations, and Applied To of a rule can contain only Antrea groups.
- Applied To is supported at both policy level and rule level. If both are specified, Applied To at the policy level takes precedence.
- Services, including raw port and protocol combination, are supported. However, the following constraints apply:
  - Only TCP and UDP services are supported. All other services are not supported.
  - In raw port and protocol combinations, TCP and UDP service types are supported.
  - Only destination ports are supported.
- Policy statistics and rule statistics are supported. Rule statistics are not aggregated for all the Antrea Kubernetes clusters to which the security policy is applied. In other words, rule statistics are displayed for each Antrea Kubernetes cluster.

Security Policy Features Not Supported for Antrea Kubernetes Clusters

- Layer 2 (Ethernet) rules based on MAC addresses are not supported.
- Layer 7 rules based on Context Profiles are not supported. For example, rules based on application ID, FQDN, and so on.
- Antrea groups with IP addresses are not supported in the Applied To of the security policy and firewall rules.
- Time-based scheduling of rules is not supported.
- Antrea groups are not supported in a firewall exclusion list. (**Security > Distributed Firewall > Actions > Exclusion List**).
- Negating or excluding the Antrea groups that you have selected in the sources or destinations of a firewall rule is not supported.
- Identity firewall is not supported.

- Global groups created for an NSX Federated environment cannot be used in security policies that are applied to Antrea Kubernetes clusters.
- Advanced policy configuration does not support the following settings:
  - TCP Strict
  - Stateful

Add a Distributed Firewall Policy to Secure Traffic Within an Antrea Kubernetes Cluster

To secure traffic between pods in an Antrea Kubernetes cluster, you can create distributed firewall policies (security policies) in NSX and apply them to one or more Antrea Kubernetes clusters.

Note This documentation uses the term "Antrea Kubernetes cluster" to mean Kubernetes clusters with Antrea CNI. The term "Kubernetes cluster" is a generic term, which represents Tanzu Kubernetes Grid (TKG) clusters with Antrea CNI, OpenShift clusters with Antrea CNI, or do it yourself (DIY) Kubernetes clusters with Antrea CNI.

The UI use the term "Antrea container cluster" for a few UI fields or labels. In the *Procedure* section of this documentation, the term "Antrea container cluster" is retained for those UI fields or labels. For all free-form text, the term "Antrea Kubernetes cluster" is used.

Prerequisites

Antrea Kubernetes clusters are registered to NSX.

Procedure

1 From your browser, log in to an NSX Manager at `https://nsx-manager-ip-address`.

2 Click the **Security** tab, and then under **Policy Management**, click **Distributed Firewall**.

The **Category Specific Rules** page is displayed.

Note NSX Manager UI fetches the information about registered Antrea Kubernetes clusters when you start the NSX Manager application in the browser. If the application UI is already open, it does not fetch the Antrea Kubernetes cluster registration information automatically. This behavior is expected and per the current UI design. If you have registered the first Antrea Kubernetes cluster after the NSX Manager application is opened, ensure that you refresh the browser after navigating to the **Category Specific Rules** page. A manual refresh ensures that the Antrea-specific UI elements are visible in the UI when you reach step 4 of this procedure.

This manual browser refresh is required only once, and not every time after a new Antrea Kubernetes cluster is registered to NSX.

3 Select the category in which you want to create the security policy.

Layer 2 (Ethernet) firewall rules based on MAC addresses are currently not supported for Antrea Kubernetes clusters. Categories in NSX correspond to Tiers in Antrea. Security policies are enforced in the Antrea Kubernetes clusters in the following descending order of precedence:

- Emergency category (highest precedence)
- Infrastructure category
- Environment category
- Application category (lowest precedence)

Within a category, firewall rules are processed top-down in the order in which the rules are set. Categories provide a means of organizing rules. For instance, multiple user roles (personas) can create security policies without overriding or conflicting the policies of each other. For example, a security administrator can create policies in the Emergency category for specific quarantine or allow rules. An application developer can create policies in the Application category to secure traffic between specific pods in an application. A network administrator can create policies in the Infrastructure category to define access rules for shared services, such as DHCP, DNS, Active Directory, and so on.

4 Click **Add Policy** and specify the policy configuration settings.

a Enter a unique name for the policy.

b By default, the policy is applied to distributed firewall. Next to **Applied To**, click the edit icon.

The **Set Applied To** page opens.

c Select the **Antrea Container Clusters** option.

d (Required) Select at least one cluster to decide the span or scope of enforcement of the security policy.

The span of the policy can either be a single Antrea Kubernetes cluster or multiple Antrea Kubernetes clusters.

e (Optional) Limit the span of the policy by selecting Antrea groups.

When you select Antrea groups in the **Applied To** of a policy, this configuration is used for all the rules in the policy. To specify a different set of Antrea groups for each rule in the policy, skip this step, and specify the **Applied To** while adding the rules in the policy.

Note Antrea groups with IP addresses must not be used in the **Applied To** of the policy because NSX cannot compute effective pod members from the IP addresses.

f (Optional) Click the gear icon at the extreme right corner to specify advanced configuration settings of the policy.

For security policies that are applied to Antrea Kubernetes clusters, **TCP Strict** and **Stateful** settings are dimmed. These settings are currently not supported.

Only **Locked** and **Comments** settings are supported. By default, a policy is not locked. To prevent multiple users from making changes to the policy, turn on the **Locked** option.

g Click **Publish**.

You can add multiple policies and then publish all of them together.

The policy status initially changes to In Progress, and after it is successfully realized in the Antrea Kubernetes clusters, the status changes to Success. If the policy realization fails due to any reason, click the `Failed` status to view the errors in a pop-up window.

5 Select the check box next to the policy name and click **Add Rule**. Enter a rule name.

By default, the **Sources**, **Destinations**, **Services**, and **Applied To** columns of the rule show `Any`.

Note Context Profiles are currently not supported for rules that are applied to Antrea Kubernetes clusters.

6 Specify the rule settings.

a In the **Sources** or **Destinations** column, click the edit icon, and select one or more Antrea groups.

The following constraints apply to specifying rule sources and destinations:

- Only Antrea groups can be selected. Groups with NSX members cannot be used in a rule. In other words, a rule cannot contain a mix of Antrea groups and groups of type **Generic**, **IP Addresses Only**.
- When groups are selected in the **Sources** column, the **Destinations** column is not applicable. You can add destinations in the **Applied To** of the rule.
- When groups are selected in the **Destinations** column, the **Sources** column is not applicable. You can add sources in the **Applied To** of the rule.

Setting **Sources** and **Applied To** in the rule filters the traffic from **Sources** to **Applied To**. Setting the **Destinations** and **Applied To** in the rule filters the traffic from **Applied To** to **Destinations**. Antrea data path does the filtering on the **Applied To** group members.

b (Optional) In the **Services** column, click the edit icon, and select the services.

If no services are selected, Any is used.

The following constraints apply to specifying services:

- Only TCP and UDP services are supported. All other services are not supported.
- Raw port and protocol combination supports only TCP and UDP service type.
- Only destination ports are supported. Source ports are not supported.

c In the **Applied To** column, click the edit icon, and select the Antrea groups to which you want to apply the rule.

If no groups are selected, Any is used.

Note

- Antrea groups with IP addresses must not be used in the **Applied To** of the rule because NSX cannot compute effective pod members from the IP addresses.
- If you specify Antrea groups in the **Applied To** of both the policy and the rule, the groups in the **Applied To** of the policy take precedence over the groups in the **Applied To** of the rule.

d From the **Action** drop-down menu, select one of these options.

Option	Description
Allow	Allows all L3 traffic with the specified source, destination, and protocol to pass through the current firewall context. Packets that match the rule, and are accepted, traverse the Kubernetes cluster as if the firewall is not present.
Drop	Drops packets with the specified source, destination, and protocol. Dropping a packet is a silent action with no notification to the source or destination. Dropping the packet causes the connection to be retried until the retry threshold is reached.
Reject	Rejects packets with the specified source, destination, and protocol. Rejecting a packet is a more graceful way to deny a packet, as it sends a destination unreachable message to the sender. If the protocol is TCP, a TCP RST message is sent. ICMP messages with administratively prohibited code are sent for UDP, ICMP, and other IP connections. A benefit of the reject action is that the sending application is notified after only one attempt that the connection cannot be established.

e Click the toggle button to turn on or turn off the rule.

By default, the rule is turned on.

f (Optional) Click the gear icon to configure other rule settings.

Rule Setting	Description
Logging	By default, logging is turned off. The firewall logs are included in the Antrea Agent logs. When you create a support bundle request for Antrea Kubernetes cluster, and select the nodes from the cluster, the support bundle includes the Antrea Agent logs for those nodes.
Direction	Refers to the direction of traffic from the perspective of the destination pod. Rule direction is read-only in the following cases, and it cannot be edited: ■ When sources are specified in the rule, the direction is In. ■ When destinations are specified in the rule, the direction is Out. Rule direction is editable when the sources and destinations are set to Any. In this case, the default direction is In-Out. However, you can change the direction to In or Out.
Comments	Enter any notes about the rule, if required. These comments are not propagated to the Antrea Kubernetes cluster. Therefore, the rule comments will not appear as annotations in the Antrea Cluster Network Policy specifications.

7 Click **Publish** to push the rules to the Antrea Kubernetes clusters.

You can add multiple rules and then publish them together.

Note After the security policy is realized in the Antrea Kubernetes clusters, you cannot edit the **Applied To** of the policy. That is, NSX does not allow you to change the span of the security policy from **Antrea Container Clusters** to **DFW** or **Groups**.

Results

The following results occur in the Antrea Kubernetes clusters:

- Antrea network plug-in creates a cluster network policy corresponding to each distributed firewall policy that is applied to the Antrea Kubernetes clusters.
- If the rules contain sources, corresponding ingress rules are created in the Antrea Cluster Network Policy.
- If the rules contain destinations, corresponding egress rules are created in the Antrea Cluster Network Policy.
- If the rules contain Any-Any configuration, Antrea Controller in the cluster splits the Any-Any rule into two rules: One ingress rule with Any to Any, and another egress rule with Any to Any.

Note Antrea network plug-in does not prevent you from updating or deleting the Antrea cluster network policies from the kubectl command line. But, you must avoid doing it. The reason is that the security policies are managed by NSX. Therefore, the Central Control Plane Adapter in the Antrea Kubernetes cluster immediately overwrites the policy changes that are made from the kubectl command line. In other words, NSX is the source of truth for the policies. The changes made to these cluster network policies through the kubectl command line are not displayed in NSX Manager.

What to do next

After the security policies are successfully realized in the Antrea Kubernetes clusters, you can do the following optional tasks:

- Verify that the Antrea cluster network policies are shown in the Kubernetes clusters. Run the following kubectl command in each Antrea Kubernetes cluster:

  ```
  $ kubectl get acnp
  ```

  Note The `priority` parameter in the Antrea cluster network policies shows a float value. This result is expected. NSX Manager UI does not display the priority of the distributed firewall policies. NSX internally assigns an integer value to the priority of each policy. This integer value is assigned from a large range. But, Antrea network plug-in assigns a smaller float number (absolute value) to the priority of Antrea cluster network policies. Therefore, the NSX priority values are internally normalized to smaller float numbers. However, the order in which you add the policies in a distributed firewall category is preserved for the Antrea cluster network policies.

  You can also view the details of the Antreacluster network policies in the NSX inventory. In NSX Manager, navigate to, **Inventory > Containers > Clusters**. Expand the cluster name and click the number next to **Cluster Network Policies** to view the details of the policies, including the YAML specifications.

- View policy statistics by using the NSX API:

  ```
  GET https://{nsx-mgr-ip}/api/v1/infra/domains{domain-id}/security-policies/{security-policy-name}/statistics?container_cluster_path=/infra/sites/{site-id}/enforcement-points/{enforcement-point-id}/cluster-control-planes/{cluster-name}
  ```

- View runtime rule statistics in the UI:

  a In NSX Manager, navigate to **Security > Distributed Firewall**.

  b Expand the policy name, and then click the graph icon at the extreme right corner of each rule.

  c Select the Kubernetes cluster from the drop-down menu to view the rule statistics for each Kubernetes cluster.

  The statistics of the rule are computed separately for each Kubernetes cluster where the rule is enforced. The statistics are not aggregated for all the Kubernetes clusters and displayed in the UI. The rule statistics are computed every minute.

Example: Add a Distributed Firewall Policy to Secure Traffic Within an Antrea Kubernetes Cluster

In this example, your goal is to create a distributed firewall policy in NSX to secure pod-to-pod traffic in the Enterprise Human Resource application, which is running in a single Antrea Kubernetes cluster.

Let us assume that the pod workloads in the Antrea Kubernetes cluster are running Web, App, and Database microservices of the Enterprise Human Resource application. You have added Antrea groups in your NSX environment by using pod-based membership criteria, as shown in the following table.

Antrea Group Name	Membership Criteria
HR-Web	Pod Tag Equals Web Scope Equals HR
HR-App	Pod Tag Equals App Scope Equals HR
HR-DB	Pod Tag Equals DB Scope Equals HR

Your objective is to create a security policy in the Application category with three firewall rules, as follows:

- Allow all traffic from HR-Web group to HR-App group.
- Allow all traffic from HR-App group to HR-DB group.
- Reject all traffic from HR-Web to HR-DB group.

Prerequisites

Antrea Kubernetes cluster is registered to NSX.

Procedure

1. From your browser, log in to an NSX Manager at `https://nsx-manager-ip-address`.
2. Click the **Security** tab, and then under **Policy Management**, click **Distributed Firewall**.

   The **Category Specific Rules** page is displayed.
3. Make sure that you are in the **Application** category.
4. Click **Add Policy** and enter a policy name.

   For example, enter **`EnterpriseHRPolicy`**.
5. In the **Applied To** of the policy, select the Antrea Kubernetes cluster where the pod workloads of the Enterprise Human Resource application are running.
6. Publish the policy.
7. Select the policy name, and click **Add Rule**.

   Configure three firewall rules, as shown in the following table.

Rule Name	Rule ID	Sources	Destinations	Services	Applied To	Action
Web-to-App	1022	HR-Web	N/A	Any	HR-App	Allow
App-to-DB	1023	HR-App	N/A	Any	HR-DB	Allow
Web-to-DB	1024	HR-Web	N/A	Any	HR-DB	Reject

The rule IDs in the table are only sample values for this example. The rule IDs can vary in your NSX environment.

8 Publish the rules.

Results

When the policy is realized successfully, the following results occur in the Antrea Kubernetes cluster:

- An Antrea cluster network policy (ACNP) is created.
- Rules 1022, 1023, and 1024 are enforced in the Kubernetes cluster in that order.
- For each firewall rule, a corresponding ingress rule is created in the cluster network policy.

Firewall Policies for Securing Traffic Between Antrea Kubernetes Clusters and VMs in an NSX Network

Starting in NSX 4.1, you can create generic groups with Kubernetes member types in dynamic membership criteria to match traffic entering into or leaving from Antrea Kubernetes clusters.

You can then use these generic groups in distributed firewall rules or gateway firewall rules to secure traffic between VMs in the NSX environment and pods in Antrea Kubernetes clusters.

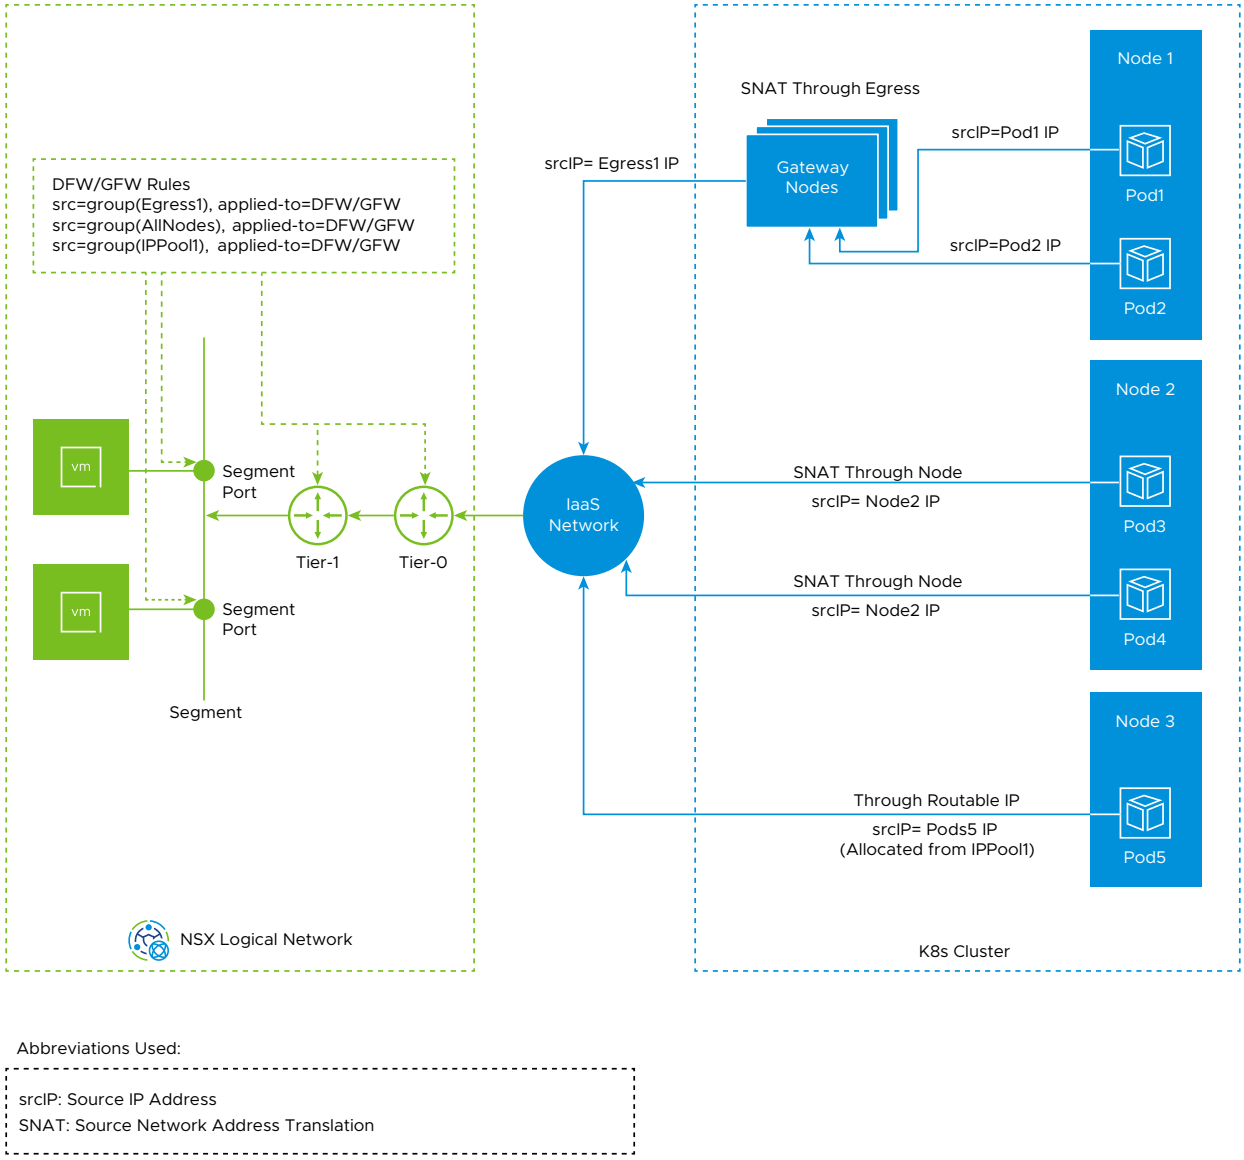
Important In a multi-tenant NSX environment, Kubernetes cluster resources are not exposed to the project inventory. Therefore, inside a project, you cannot create generic groups with Kubernetes member types in dynamic membership criteria. You must create the groups in the **Default** view (default space) of NSX, and then consume them in the distributed firewall rules or gateway firewall rules under the default space.

Securing Traffic from Antrea Kubernetes Clusters to VMs in NSX

To match traffic from Antrea Kubernetes clusters to NSX, you can add generic groups based on the following Kubernetes member types (resources) in the firewall rules:

- Kubernetes Node
- Antrea Egress
- Antrea IP Pool

The following diagram shows the typical traffic flows from the Antrea Kubernetes cluster to the VMs in the NSX logical network.



Abbreviations Used:

srcIP: Source IP Address
SNAT: Source Network Address Translation

As shown in this traffic flow diagram, to match traffic that is leaving the Antrea Kubernetes cluster, the following scenarios are possible:

Scenario 1: SNAT through Egress (gateway nodes)

You can deploy an Egress resource in your Antrea Kubernetes cluster. In the manifest file of the Egress resource, select the grouping criteria of pods to which the Egress resource applies to. Either you can specify the egress IP manually in the manifest file or Antrea CNI can allocate an egress IP from an ExternalIPPool custom resource.

Important The egress IP addresses must be routable on the physical network and they must be reachable from all the nodes in the K8s cluster.

Antrea CNI sends the outgoing traffic from the selected pods to the gateway nodes. The gateway nodes perform a source address translation (SNAT) to translate the source IP (pod IP) to the egress IP.

For example, in the preceding traffic flow diagram, outgoing traffic from pod1 and pod2, which are running on node1, is sent to the gateway nodes. The gateway node translates the pod IP addresses to egress1 IP.

To learn more about Egress resource, see the Egress Guide on the Antrea documentation portal.

The cluster nodes that can be used as gateway nodes depend on the node network topology, which is controlled by your physical network administrator or the IaaS network administrator. The gateway nodes have the following two characteristics:

- The physical network allows the nodes to have a routable IP as a source IP to access the external network.
- Optionally, a physical firewall on the physical gateway is configured to filter outgoing traffic from the gateway nodes.

Consider the following examples:

- Examp1e 1: Assume that your Antrea Kubernetes cluster has three nodes: node1, node2, and node3. Let us say that node1 and node2 in the cluster are special nodes because they are assigned two network interfaces. One interface for K8s cluster communication and another interface for connecting to the external network. Node3 is assigned only one network interface for cluster communication. In this scenario, only node1 and node2 can be used as gateway nodes.
- Example 2: Assume that all three nodes in the K8s cluster are assigned only one network interface, and the network administrator has configured the physical firewall to allow only node-1 and node-2 MAC addresses to visit the Internet. In this scenario, node1 and node2 can be used as gateway nodes.
- Example 3: Assume that all the three nodes can visit the Internet. The network administrator has configured a physical firewall on the physical gateway, and the firewall wants to filter traffic from a specific egress IP. In this scenario, all three nodes can be used as gateway nodes.

The following paragraph explains an implementation-specific limitation:

Any physical host or a VM that is outside the Antrea Kubernetes cluster cannot be used as a gateway node for tunnelling the pod's outgoing traffic. The reason is that Antrea Agent maintains the egress IP addresses, and this agent runs as a pod on each node of the K8s cluster. Therefore, the physical host or a VM that you want to use as a gateway node must be a part of the K8s cluster.

Scenario 2: SNAT through nodes

In this scenario, pods are not explicitly selected in the manifest file of any Egress resource.

The pod's outgoing traffic is masqueraded to the node IP, and then the traffic is sent out to the IaaS network. For example, in the preceding traffic flow diagram, an SNAT rule translates the source IP (pod3 IP, pod4 IP) to the node2 IP, and then the pod's outgoing traffic is sent out to the IaaS network.

Scenario 3: Routable pods bridged directly to the IaaS network

Routable pods mean that pods are assigned routable IP addresses on the underlay network. Usually, pods are assigned private IP addresses from the PodCIDR property, which is configured on the node. When pods want to access a network that is external to the K8s cluster, the traffic requires an SNAT operation to translate the source IP (pod IP) to a routable IP.

If your Antrea Kubernetes cluster is deployed on the NSX overlay network with the Tanzu Kubernetes Grid Service (TKGS), then there is a mechanism for the Tanzu Kubernetes cluster (TKC) nodes to request routable subnets from NSX, and use this routable subnet as the node's PodCIDR property. Then, the pods can get routable IP addresses. To learn more about configuring a cluster with a routable pods network, see Example: TKC with Routable Pods Network in the *VMware vSphere 8.0* documentation.

If your Antrea Kubernetes cluster is not deployed on the NSX overlay network, the Kubernetes administrator can ensure that a routable subnet or a routable IP range is reserved in the underlay network. The administrator can configure the subnet or an IP range in the manifest file of the IPPool custom resource. Then, Antrea CNI can allocate routable IP addresses to the pods.

For example, in the preceding traffic flow diagram, pod5 that is running on node5 gets a routable IP from an IPPool custom resource.

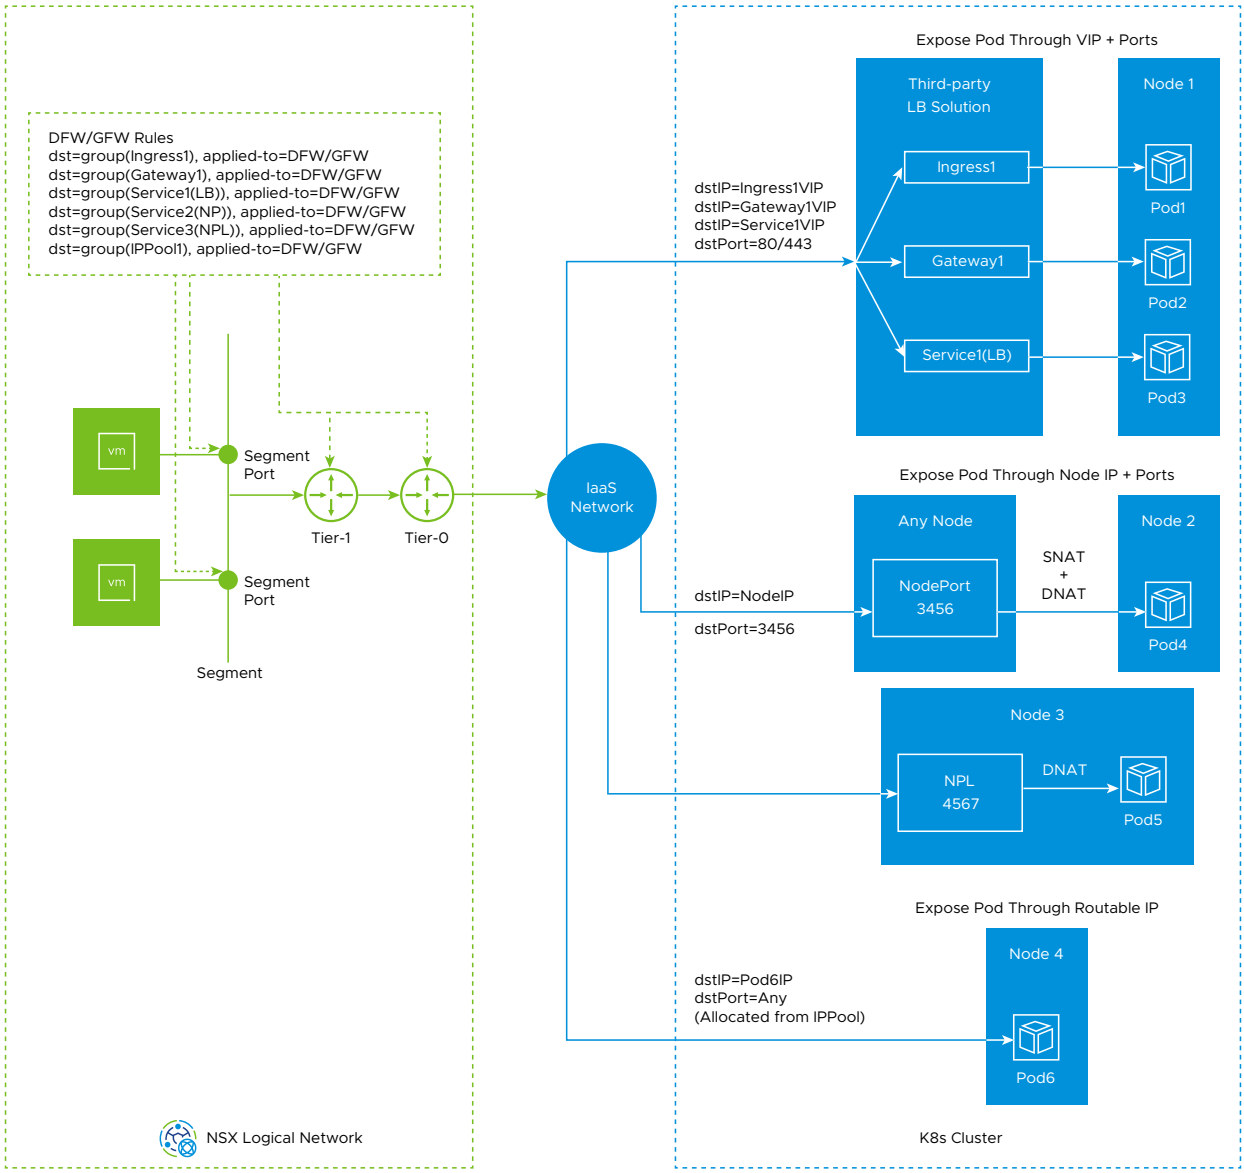
To learn more about IPPool custom resource, see the Antrea IPAM Guide on the Antrea documentation portal.

Securing Traffic from VMs in NSX to Antrea Kubernetes Clusters

To match traffic from NSX to Antrea Kubernetes Clusters, you can add generic groups based on the following Kubernetes member types (resources) in the firewall rules:

- Kubernetes Service
- Kubernetes Ingress
- Kubernetes Gateway
- Antrea IP Pool

The following diagram shows the typical traffic flows from VMs in the NSX logical network to the Antrea Kubernetes cluster.



Abbreviations Used:

dstIP: Destination IP Address
dstPort: Destination Port
VIP: Virtual IP Address
SNAT: Source Network Address Translation
DNAT: Destination Network Address Translation
NPL: NodePortlLocal

As shown in this traffic flow diagram, to match traffic that is entering into the Antrea Kubernetes cluster, the following scenarios are possible:

Scenario 1: Expose pods to external traffic through virtual IP and ports

You can expose pods to external traffic by implementing a third-party load balancer solution with these native Kubernetes resources:

- Ingress: It functions as a layer 7 load balancer. An Ingress resource will provide a virtual IP address (VIP) and expose some ports (TCP/UDP).
- Gateway: It functions as a layer 4 and layer 7 load balancer. A Gateway resource will provide a VIP and expose some ports (TCP/UDP).
- Service of type LoadBalancer: It functions as a layer 4 load balancer. The service will provide a VIP and expose some ports (TCP/UDP).

Note Although Ingress and Gateway resources can function as a layer 7 load balancer, for distributed firewall rules and gateway firewall rules in NSX, they appear as layer 3 or layer 4 load balancer by using IP addresses and ports to match the incoming traffic.

Scenario 2: Expose pods to external traffic through node IP and ports

- Kubernetes service of type NodePort: It triggers all the nodes in the K8s cluster to open a port on the node for each service port. The service becomes reachable through node IP and port. The node performs SNAT and DNAT on the traffic.
- Kubernetes service of type ClusterIP with NodePortLocal feature enabled: It requires you to enable the NodePortLocal feature on the Antrea Agent and add annotation in the manifest file of the Kubernetes service. Antrea CNI recognizes the annotation and opens a port for each pod on the node where the pod is running. NodePortlLocal avoids opening a port on all the nodes of the K8s cluster, thereby saving available port ranges. It also avoids an SNAT operation, and preserves the original client IP.

  To learn more about NodePortLocal feature, see the NodePortLocal Guide on the Antrea documentation portal.

Scenario 3: Expose pods to external traffic through routable pod IP addresses

Antrea supports assigning routable IP addresses to pods by deploying the IPPool custom resource in your Antrea Kubernetes cluster. The pods are allocated IP address from this pool and they are bridged directly to the IaaS network.

Supported IaaS Networks

Antrea Kubernetes clusters can be deployed on the following IaaS platforms:

- vSphere-based on-premises data center: as Tanzu Kubernetes clusters, which are created by using the Tanzu Kubernetes Grid Service. The clusters are managed by Tanzu Kubernetes Grid (TKG) 2.0 and vSphere.
- VMware Cloud: as Tanzu Kubernetes clusters that are managed by TKG 2.0 and VMC SDDC.
- Public clouds: as managed Kubernetes clusters on the public cloud platforms.
- Physical servers: as Kubernetes clusters on bare-metal servers.

- OpenShift:
  - Antrea CNI deployed in OpenShift clusters, and OpenShift is deployed in installer provisioned infrastruture (IPI) mode or user provisioned infrastructure (UPI) mode.
  - Infrastructure providers can be any one of the following: vSphere, Amazon Web Services (AWS), Azure, OpenStack, Google Cloud Platform (GCP), bare metal.

Important The IaaS network is responsible for traffic between Antrea Kubernetes clusters and VMs. The traffic between different sites, such as on-premises data centers, VMC, and public cloud might involve an IaaS-specific VPN. In all cases, there might be SNAT operations applied by the IaaS network. The NSX administrator must ensure that the source or destination IP addresses are routable so that they can be used in firewall rules to protect traffic between VMs and Kubernetes clusters.

Recommended Approaches to Apply Firewall Policies

Both distributed firewall and gateway firewall allow you to specify generic groups with Kubernetes member types in the sources and destinations of the firewall rules. So, you have the flexibility to decide whether to reference the groups in the distributed firewall or gateway firewall.

VMware provides the following recommendation:

- Use NSX Gateway Firewall to allow or block traffic from Antrea Kubernetes clusters to VMs in an NSX network.

  This approach helps you to filter the traffic at the earliest point when traffic enters into the NSX overlay network.

- Use NSX Distributed Firewall to allow or block traffic from VMs in an NSX network to Antrea Kubernetes clusters.

  This approach helps you to filter the traffic at the earliest point when traffic leaves the VMs in an NSX overlay network.

Summary of Kubernetes Member Types for Usage in Firewall Policies

The following table summarizes the Kubernetes member types that you can use in NSX generic groups to match traffic in firewall rules.

Member Type	Scope	Usage in Firewall Policy
Kubernetes Cluster	Cluster	Appears as an AND condition in the dynamic group criteria to match resources from specific clusters.
Kubernetes Namespace	Namespace	Appears as an AND condition in the dynamic group criteria to match resources from specific namespaces.
Antrea Egress	Cluster	Match outgoing traffic from Antrea Kubernetes clusters where source IP = egress IP

Member Type	Scope	Usage in Firewall Policy
Antrea IP Pool	Cluster	Match outgoing traffic from Antrea Kubernetes clusters where source IP is in IP ranges. Match traffic entering into Antrea Kubernetes clusters where destination IP is in IP ranges.
Kubernetes Node	Cluster	Match outgoing traffic from Antrea Kubernetes clusters where source IP is among the cluster node IP addresses.
Kubernetes Ingress	Namespace	Match traffic entering into Antrea Kubernetes clusters where destination IP is in Ingress IP addresses.
Kubernetes Gateway	Namespace	Match traffic entering into Antrea Kubernetes clusters where destination IP is in Gateway IP addresses.
Kubernetes Service (type=LoadBalancer)	Namespace	Match traffic entering into Antrea Kubernetes clusters where destination IP is in LoadBalancer IP addresses.
Kubernetes Service (type=NodePort)	Namespace	Match traffic entering into Antrea Kubernetes clusters where destination IP + port is in node IP addresses + NodePort range.
Kubernetes Service (type=ClusterIP)	Namespace	None
Kubernetes Service (type=ClusterIP) and NodePortLocal feature enabled	Namespace	Match traffic entering into Antrea Kubernetes clusters where destination IP + port is in node IP addresses + NodePortLocal range.

Deregister an Antrea Kubernetes Cluster from NSX

Use the `deregisterjob.yaml` file that is included with the `antrea-interworking-`*`version`*`.zip` file to deregister an Antrea Kubernetes cluster from NSX.

Prerequisites

Open the `deregisterjob.yaml` file in a text editor and replace the image URLs with any one these container images that are hosted on VMware Harbor Registry.

- `projects.registry.vmware.com/antreainterworking/interworking-debian:`*`version`*
- `projects.registry.vmware.com/antreainterworking/interworking-ubuntu:`*`version`*
- `projects.registry.vmware.com/antreainterworking/interworking-photon:`*`version`*

For *version* information, see the *VMware Container Networking with Antrea* release notes at https://docs.vmware.com/en/VMware-Container-Networking-with-Antrea/index.html.

If your Kubernetes cluster does not have Internet access, you can find the container images in the `antrea-interworking-`*`version`*`.zip` file, which you downloaded before registering the Kubernetes cluster. If the interworking deployment is running in the Kubernetes cluster, the container images are already loaded on the cluster nodes. The deregister job and the interworking deployment use the same container images.

Procedure

1 Run the following kubectl command to submit the `deregisterjob.yaml` file to the Kubernetes API server.

```
$ kubectl apply -f deregisterjob.yaml
```

This job takes some time to complete. The following actions happen in the background:

- Deletes the interworking deployment.
- Deletes the resources of the Antrea Kubernetes cluster from the NSX inventory. The resources include pods, services, ingress rules, nodes, and network policies.
- Removes references to the Antrea Kubernetes cluster in the distributed firewall policies and groups (if any).
- Deletes the Antrea custom resources, which were managed by NSX from the Kubernetes cluster. These custom resources include: Traceflow, ClusterNetworkPolicy, ClusterGroup, and Tier.

Note If you deregister an Antrea Kubernetes cluster by using the following NSX API, the resources of the Kubernetes cluster from the NSX inventory are deleted, and the references to the Kubernetes cluster in the NSX-managed distributed firewall policies and groups are deleted:

```
DELETE https://{nsx-mgr-ip}/policy/api/v1/infra/sites/{siteid}/enforcement-points/
{enforcementpoint-id}/cluster-control-planes/{cluster-name}?cascade=true
```

For more information about this API, see the *NSX API Guide*. But this API does not delete the Kubernetes cluster from the NSX inventory. You can run a curl command from the command line to delete the cluster. For more information, see Delete an Antrea Kubernetes Cluster from the NSX Inventory by Using the Command Line.

2 Check the status of the deregister job by running the following kubectl command:

```
kubectl get job -o wide deregister -n vmware-system-antrea
```

Wait for the job to complete. You might have to run this command a few times to check the status of the job.

3 After the deregister job is completed, run the following kubectl command to delete the `vmware-system-antrea` namespace and the role-based access control (RBAC) resources.

```
kubectl delete -f interworking.yaml --ignore-not-found
```

The `ignore-not-found` flag is used in this command to avoid the `Resource Not Found` error in the command output if some resources are not found for deletion.

RBAC resources, such as ServiceAccount, ClusterRole, and ClusterRoleBinding are deleted.

What to do next

Verify that the Antrea Kubernetes cluster is not shown in the NSX inventory.

1 In the NSX Manager UI, navigate to **Inventory > Containers > Clusters**.

2 Observe that the Antrea Kubernetes cluster is not shown in the inventory.

Optional: After deregistering the Antrea Kubernetes cluster, delete the principal identity (PI) user and the self-signed certificate.

- To delete the PI user account, navigate to **System > User Management > User Role Assignment**. Next to the PI user name, click  , and then click **Delete**.
- To delete the self-signed certificate, navigate to **System > Certificates**. Next to the certificate name, click  , and then click **Delete**.

If required, you can re-register the same Antrea Kubernetes cluster. However, before re-registering the cluster, ensure that you have run the following kubectl command:

```
kubectl delete -f interworking.yaml --ignore-not-found
```

To re-register the same Antrea Kubernetes cluster, do any one of the following:

- If you want to reuse the same PI user account and self-signed certificate for re-registering the Antrea Kubernetes cluster, do not delete PI user account and the self-signed certificate from NSX. In this case, no changes are required in the `bootstrap-config.yaml`.
- If you want to use a new PI user account and self-signed certificate for re-registering the Antrea Kubernetes cluster, delete the old PI user account and the self-signed certificate. Begin the process by creating a self-signed certificate and use this new certificate to create a PI user account. Edit the `tls.crt` and `tsl.key` arguments in the `bootstrap-config.yaml` with the information of this new PI user.

Delete an Antrea Kubernetes Cluster from the NSX Inventory by Using the Command Line

After you have deregistered an Antrea Kubernetes cluster by running an NSX API, follow the instructions in this documentation to delete the Kubernetes cluster from the NSX inventory by using the command line.

Prerequisites

- You must have the user name and password of the NSX Enterprise Admin user.
- You must be able to connect to NSX Manager UI and API.

Procedure

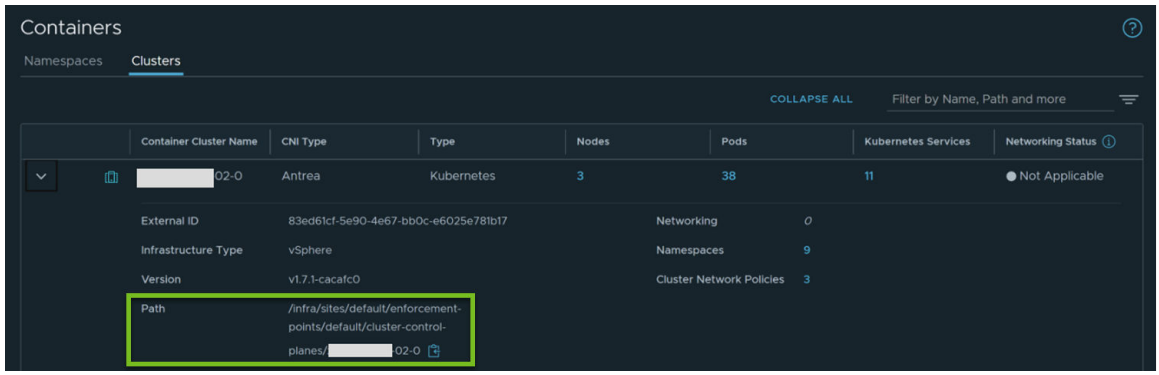
1 Find the path of the Antrea Kubernetes cluster in the NSX Manager UI.

a From your browser, log in to an NSX Manager at https://*nsx-manager-ip-address*.

b Navigate to **Inventory > Containers > Clusters**.

c Expand the cluster to be deleted and copy the text that you see next to the **Path** field.

For example:



2 Run the following curl command from the command line to delete the Antrea Kubernetes cluster:

```
curl -k -u '{AdminUserName}:{AdminPassword}' \
 -X DELETE -H "X-Allow-Overwrite: true" \
 https://{NSX-Mgr-IP}/policy/api/v1{Path}?cascade=true
```

In this command:

- Replace *{AdminUserName}*, *{AdminPassword}*, and *{NSX-Mgr-IP}* with their actual values as applicable to your NSX environment.
- Replace *{Path}* with the text that you copied in step 1.

For example:

```
curl -k -u 'Admin:Password123' \
 -X DELETE -H "X-Allow-Overwrite: true" \
 https://192.168.1.1/policy/api/v1/infra/sites/default/enforcement-points/default/cluster-control-planes/cluster-sales?cascade=true
```

Troubleshooting Antrea to NSX Integration Issues

Use this section to learn about collecting logs from the Antrea Kubernetes cluster nodes and resolving problems that you might encounter when integrating Antrea Kubernetes clusters to NSX.

Collect Support Bundles for an Antrea Kubernetes Cluster

You can use the Support Bundle feature in NSX to collect log files from Antrea Kubernetes cluster nodes for detailed troubleshooting and diagnostic purposes.

NSX can upload the support bundles to an NSX Manager node from where you triggered the support bundle collection request, or it can upload the support bundles to a remote file server that you specified in the request. If the support bundles are uploaded to an NSX Manager node, you can download them to your local computer.

Note This documentation uses the term "Antrea Kubernetes cluster" to mean Kubernetes clusters with Antrea CNI. The term "Kubernetes cluster" is a generic term, which represents Tanzu Kubernetes Grid (TKG) clusters with Antrea CNI, OpenShift clusters with Antrea CNI, or do it yourself (DIY) Kubernetes clusters with Antrea CNI.

The UI use the term "Antrea container cluster" for a few UI fields or labels. In the *Procedure* section of this documentation, the term "Antrea container cluster" is retained for those UI fields or labels. For all free-form text, the term "Antrea Kubernetes cluster" is used.

A support bundle for an Antrea Kubernetes cluster contains log files for the following components:

- Antrea Controller
- Antrea Agent
- Management Plane Adapter
- Central Control Plane Adapter
- Open vSwitch

Supported and Unsupported Features

- From an NSX Manager node, you can start only a single support bundle collection request. But, you can collect support bundles for multiple Antrea Kubernetes clusters with a single collection request.
- If you are using an NSX Manager cluster with three Management nodes, you can start a separate support bundle collection request simultaneously from each NSX Manager node. However, the Antrea Kubernetes cluster that you select in each collection request must be different.

  For example, assume that you have started a support bundle collection request on NSX Manager node A. In this collection request, you selected cluster nodes 1 and 2 from Antrea Kubernetes cluster X. Simultaneously, if you start a second collection request on NSX Manager node B for the cluster nodes 3 and 4 in the Antrea Kubernetes cluster X, one of these two collection requests will fail. You must wait for the first collection request to complete before triggering the second request for the same Antrea Kubernetes cluster.

- Collection of support bundles from NSX Manager Central CLI is currently not supported for Antrea Kubernetes clusters.

Prerequisites

Antrea Kubernetes clusters are registered to NSX.

Procedure

1 From your browser, log in to an NSX Manager at `https://nsx-manager-ip-address`.

2 Navigate to **System > Support Bundle**.

The **Request Bundle** page opens.

Important NSX Manager UI fetches the information about registered Antrea Kubernetes clusters when you start the NSX Manager application in the browser. If the application UI is already open, it does not fetch the Antrea Kubernetes cluster registration information automatically. This behavior is expected and per the current UI design. If you have registered the first Antrea Kubernetes cluster after the NSX Manager application is opened, ensure that you refresh the browser after navigating to the **Request Bundle** page. A manual refresh ensures that you can select Antrea Kubernetes clusters as the target nodes in the next step of this procedure.

This manual browser refresh is required only once, and not every time after a new Antrea Kubernetes cluster is registered to NSX.

3 Select the target nodes to include in the support bundle request.

The available types of nodes are:

- Antrea Container Clusters
- Management Nodes
- Edges
- Hosts
- Public Cloud Gateways

A single support bundle request can include a mix of different types of nodes in the NSX environment. For example, you can select nodes from Antrea Kubernetes clusters, NSX Manager nodes, and NSX Edge nodes in the same collection request. However, the scope of this procedure is to explain the workflow of creating a support bundle collection request for only Antrea Kubernetes clusters.

4 From the **Type** drop-down menu, select **Antrea Container Clusters**.

5 From the **Container Cluster** list, select the name of a cluster.

   If the list has several clusters to select from, enter the first few characters of the cluster name. System filters the list and displays only the cluster names that match the characters you have entered.

   All nodes in the selected cluster are displayed in the **Available** list.

6 Select one or multiple nodes from the cluster and click the right arrow to move them to the **Selected** list.

   To select nodes from multiple clusters in a single collection request, repeat steps 4 and 5 for each cluster.

7 (Optional) In the **Log age (days)** text box, keep the default value or enter the specific number of days' worth of logs that you want the support bundle to include. Specify the log age as a number of days.

8 (Optional) To upload the support bundle to a remote file server, specify the file server settings.

   a Enter an IP address or the host name of the remote file server.

   b Enter the file transfer protocol and port number. Default port number is 22.

   c Enter the user name and password to access the remote file server.

   d Enter the path to the destination folder where the support bundle file is to be uploaded.

   When remote file server settings are not specified, the support bundle is uploaded to the NSX Manager node from where you triggered the support bundle collection request.

9 Click **Start Bundle Collection**.

   The runtime details of the collection request are displayed on the **Status** page. The collection process takes a few minutes. The time taken to create the support bundle depends on the number of log files to collect from each node in the container cluster.

10 After the collection process is complete, click **Download**.

   The support bundle file is saved on your local computer. If you had specified remote file server settings, the **Download** button is not displayed in the UI.

Results

A support bundle collection request generates a single tape archive (TAR) file with the following file naming convention: `nsx_support_archive_datestamp_timestamp.tar`

Support bundle collection request can fail in the following situations:

- If the Antrea NSX Adapter on a Kubernetes cluster fails when the support bundle request is in progress, the collection of logs fails for that Kubernetes cluster.

- If the NSX Manager Appliance fails or is not reachable when the support bundle request is in progress, the collection of logs fails. Until the connectivity issue to the NSX Manager is resolved, you can use the native command line tool of Antrea (antctl) to collect log files from the Antrea Kubernetes clusters.

Partial Success Scenario

Consider that you selected 10 nodes from a single Antrea Kubernetes cluster for the support bundle collection. During the collection process, log files were collected successfully from five nodes in the Antrea Kubernetes cluster, but were not collected for the remaining five nodes. In other words, the collection request succeeded partially. In this situation, the collection request status is `Successful` and the support bundle file (TAR) contains logs for the five successful nodes.

What to do next

1 Extract the TAR file. The following files are displayed.

File Name	Description
`manifest.json`	This file contains a summary of the collection request results and the properties of the collection request. For example, it contains information about: ■ The nodes for which the collection succeeded. ■ The nodes for which the collection failed. ■ The cluster IDs and node IDs that were used in the collection request.
`nsx_antrea_`*`cluster-id`*`.tgz`	A single `.tgz` archive file is created for each Antrea Kubernetes cluster in the support bundle.

2 Extract the `nsx_antrea_`*`cluster-id`*`.tgz` file. The following files are displayed.

File Name	Description
`adapters.tar.gz`	This archive file contains the log files of the Management Plane Adapter and the Central Control Plane Adapter.
`agent_`*`node_name`*`.tar.gz`	This archive file contains the log files of the Antrea Agent and Open vSwitch. One archive file is generated for each Antrea Kubernetes cluster node in the collection request. On extracting this archive file, you can view the following files: ■ `agentinfo` file ■ Agent logs at `/logs/agent` ■ Open vSwitch logs at `/logs/ovs` ■ OpenFlow dump ■ IPtables ■ Route dump

File Name	Description
`clusterinfo`	This file is generated for each Antrea Kubernetes cluster in the support bundle request. The file contains information about the following Kubernetes resources that are collected from the Kubernetes API server: ■ Pods ■ Nodes ■ Deployments ■ ReplicaSets ■ DaemonSets
`controller.tar.gz`	This archive file contains the log files of the Antrea Controller. On extracting this archive file, you can view the following files: ■ `controllerinfo` file ■ Controller logs at `/logs/controller`

Antrea Kubernetes Cluster Status is Down

If the status of the Antrea Kubernetes cluster is down, follow the steps in this documentation either to determine the cause of this issue and recover from it, or collect the support bundle.

Problem

The cluster control plane node is down. The Antrea Kubernetes cluster is disconnected from the Central Control Plane (CCP).

Cause

In the NSX Manager UI, navigate to **System > Fabric > Nodes > Container Clusters > Antrea**. If required, filter the list of clusters on the **Antrea** page with the **External ID** field.

Click the **Status** column of the problematic cluster. If all the components are down, the possible causes are:

- The Kubernetes cluster is deleted.
- Network connectivity issue with the CCP.
- The adapters are crashed or deleted for some reason.
- The client certificate of the adapters is incorrect.
- The version of the adapters is incompatible with the CCP.

If only the Central Control Plane Adapter is down, the CCP Adapter might have crashed.

Solution

1 If the Kubernetes cluster is deleted, clean up the leftover registration and inventory data in NSX. See Delete an Antrea Kubernetes Cluster from the NSX Inventory by Using the Command Line.

2 Get the kubectl and kubeconfig access for the Kubernetes cluster. Use kubectl to retrieve the node name on which the interworking pod is running. Start an SSH session to the node and use the curl or nc command to connect to every NSX Manager IP on ports 1234 and 1235. If the connection cannot be established, the cause is network connectivity issue with the CCP.

   Example of the curl command:

   Ensure that you replace *NSX-Manager-IP* with the IP address of NSX Manager in your environment.

```
curl -v NSX-Manager-IP:1235

Trying NSX-Manager-IP...
Connected to NSX-Manager-IP (NSX-Manager-IP) port 1235 (#0)
...
Empty reply from server
Connection #0 to host NSX-Manager-IP left intact
curl: (52) Empty reply from server
```

   Example of the nc command:

```
nc -v NSX-Manager-IP 1235 < /dev/null

Ncat: Version 7.50 (https://nmap.org/ncat)
Ncat: Connected to NSX-Manager-IP:1235.
Ncat: 0 bytes sent, 0 bytes received in 0.37 seconds.
```

3 Use kubectl to check whether all containers of the interworking pod in the vmware-system-antrea namespace are up.

   If any container is down, use kubectl to get logs of the crashed containers and check the error message. This step can help you identify failure due to any of these reasons:

   - The adapters are crashed or deleted for some reason.
   - CCP Adapter is crashed.

   Example of the kubectl command for getting the interworking pod:

```
kubectl get pod -o wide -l app=antrea-interworking -n vmware-system-antrea
```

   Note down the interworking pod name.

   Example of the kubectl command for getting the detailed state of the interworking pod:

   Ensure that you replace *pod-name* with the actual pod name.

```
kubectl get pod -o yaml pod-name -n vmware-system-antrea
```

   Example of the kubectl command for getting container logs:

Ensure that you replace *pod-name* with the actual pod name.

```
kubectl logs pod-name -c mp-adapter -n vmware-system-antrea > mp-adapter.log
kubectl logs pod-name -c ccp-adapter -n vmware-system-antrea > ccp-adapter.log
kubectl logs pod-name -c tn-proxy -n vmware-system-antrea > tn-proxy.log
kubectl logs pod-name -c election-runner -n vmware-system-antrea > election-runner.log
```

If the vmware-system-antrea namespace is missing or the interworking pod is missing, the adapters might have been deleted from the Kubernetes cluster without running the deregistration steps. You can clean up the leftover registration data and inventory from the system, and then register the Kubernetes cluster again. The cluster ID will be different after reregistering the cluster. If there is any Antrea policy applied to the cluster, you must apply the policy again after reregistering the cluster.

For instructions about cleaning up the leftover registration data, see Delete an Antrea Kubernetes Cluster from the NSX Inventory by Using the Command Line.

For instructions about registering an Antrea Kubernetes cluster to NSX, see Registering an Antrea Kubernetes Cluster to NSX.

4 Use kubectl to get nsx-proxy container logs from the interworking pod, and check the error messages.

This step can help you identify failure due to any of these reasons:

- The client certificate of the adapters is incorrect.
- The version of the adapters is incompatible with the CCP.

For example commands, see step 3.

5 If the Management Plane Adapter is up, use the support bundle feature in NSX to collect log files for the Antrea Kubernetes cluster.

For more information, see Collect Support Bundles for an Antrea Kubernetes Cluster.

Configuring NSX in Manager Mode

25

NSX has two user interface modes: Policy mode and Manager mode. If you have objects that were created in Manager mode, you should continue to use Manager mode to make changes.

For more information about the two modes, see Chapter 1 NSX Manager.

If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

This chapter includes the following topics:

- Logical Switches in Manager Mode
- Logical Routers in Manager Mode
- NAT in Manager Mode
- Grouping Objects in Manager Mode
- DHCP in Manager Mode
- IP Address Management in Manager Mode
- Load Balancing in Manager Mode
- Firewall in Manager Mode

Logical Switches in Manager Mode

You can configure logical switches and related objects in **Manager** mode. A logical switch reproduces switching functionality, broadcast, unknown unicast, multicast (BUM) traffic, in a virtual environment decoupled from the underlying hardware.

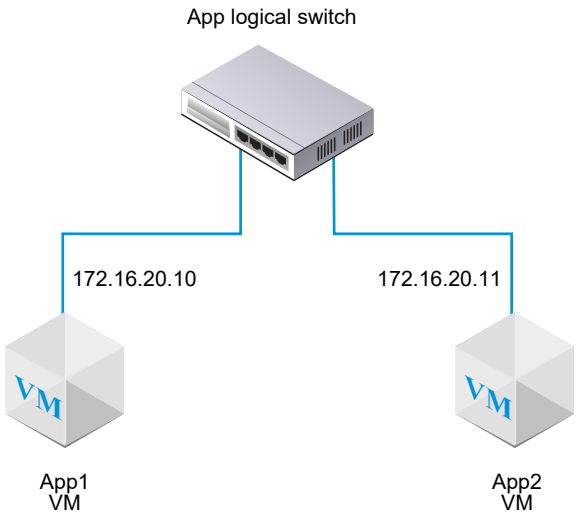
Note If you use **Manager** mode to modify objects created in the **Policy** mode, some settings might not be configurable. These read-only settings have this icon next to them: ⊖. See Chapter 1 NSX Manager for more information.

Logical switches are similar to VLANs, in that they provide network connections to which you can attach virtual machines. The VMs can then communicate with each other over tunnels between hypervisors if the VMs are connected to the same logical switch. Each logical switch has a virtual network identifier (VNI), like a VLAN ID. Unlike VLAN, VNIs scale well beyond the limits of VLAN IDs.

To see and edit the VNI pool of values, log in to NSX Manager, navigate to **Fabric > Profiles**, and click the **Configuration** tab. Note that if you make the pool too small, creating a logical switch will fail if all the VNI values are in use. If you delete a logical switch, the VNI value will be re-used, but only after 6 hours.

When you add logical switches, it is important that you map out the topology that you are building.

Figure 25-1. Logical Switch Topology



For example, the topology above shows a single logical switch connected to two VMs. The two VMs can be on different hosts or the same host, in different host clusters or in the same host cluster. Because the VMs in the example are on the same virtual network, the underlying IP addresses configured on the VMs must be in the same subnet.

NSX Cloud Note If using NSX Cloud, see NSX Features Supported with NSX Cloud for a list of auto-generated logical entities, supported features, and configurations required for NSX Cloud.

Understanding BUM Frame Replication Modes

Each host transport node is a tunnel endpoint. Each tunnel endpoint has an IP address. These IP addresses can be in the same subnet or in different subnets, depending on your configuration of IP pools or DHCP for your transport nodes.

When two VMs on different hosts communicate directly, unicast-encapsulated traffic is exchanged between the two tunnel endpoint IP addresses associated with the two hypervisors without any need for flooding.

However, as with any Layer 2 network, sometimes traffic that is originated by a VM needs to be flooded, meaning that it needs to be sent to all of the other VMs belonging to the same logical switch. This is the case with Layer 2 broadcast, unknown unicast, and multicast traffic (BUM traffic). Recall that a single NSX logical switch can span multiple hypervisors. BUM traffic originated by a VM on a given hypervisor needs to be replicated to remote hypervisors that host other VMs that are connected to the same logical switch. To enable this flooding, NSX supports two different replication modes:

- Hierarchical two-tier (sometimes called MTEP)
- Head (sometimes called source)

Hierarchical two-tier replication mode is illustrated by the following example. Say you have Host A, which has VMs connected to virtual network identifiers (VNIs) 5000, 5001, and 5002. Think of VNIs as being similar to VLANs, but each logical switch has a single VNI associated with it. For this reason, sometimes the terms VNI and logical switch are used interchangeably. When we say a host is on a VNI, we mean that it has VMs that are connected to a logical switch with that VNI.

A tunnel endpoint table shows the host-VNI connections. Host A examines the tunnel endpoint table for VNI 5000 and determines the tunnel endpoint IP addresses for other hosts on VNI 5000.

Some of these VNI connections will be on the same IP subnet, also called an IP segment, as the tunnel endpoint on Host A. For each of these, Host A creates a separate copy of every BUM frame and sends the copy directly to each host.

Other hosts' tunnel endpoints are on different subnets or IP segments. For each segment where there is more than one tunnel endpoint, Host A nominates one of these endpoints to be the replicator.

The replicator receives from Host A one copy of each BUM frame for VNI 5000. This copy is flagged as Replicate locally in the encapsulation header. Host A does not send copies to the other hosts in the same IP segment as the replicator. It becomes the responsibility of the replicator to create a copy of the BUM frame for each host it knows about that is on VNI 5000 and in the same IP segment as that replicator host.

The process is replicated for VNI 5001 and 5002. The list of tunnel endpoints and the resulting replicators might be different for different VNIs.

With head replication also known as headend replication, there are no replicators. Host A simply creates a copy of each BUM frame for each tunnel endpoint it knows about on VNI 5000 and sends it.

If all the host tunnel endpoints are on the same subnet, the choice of replication mode does not make any difference because the behaviour will not differ. If the host tunnel endpoints are on different subnets, hierarchical two-tier replication helps distribute the load among multiple hosts. Hierarchical two-tier is the default mode.

Create a Logical Switch in Manager Mode

Logical switches attach to single or multiple VMs in the network. The VMs connected to a logical switch can communicate with each other using the tunnels between hypervisors.

Note that there is a delay of up to 4 minutes in creating the first logical switch if you enable the lockdown mode.

Prerequisites

- Verify that a transport zone is configured. See the *NSX Installation Guide*.
- Verify that fabric nodes are successfully connected to NSX management plane agent (MPA) and NSX local control plane (LCP).

  In the `GET https://<nsx-mgr>/api/v1/transport-nodes/<transport-node-id>/state` API call, the `state` must be `success`. See the *NSX Installation Guide*.
- Verify that transport nodes are added to the transport zone. See the *NSX Installation Guide*.
- Verify that the hypervisors are added to the NSX fabric and VMs are hosted on these hypervisors.
- Familiarize yourself with the logical switch topology and BUM frame replication concepts. See Logical Switches in Manager Mode and Understanding BUM Frame Replication Modes.
- Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Select **Networking > Logical Switches > Switches > Add**.
3. Enter a name for the logical switch and optionally a description.
4. Select a transport zone for the logical switch.

   VMs that are attached to logical switches that are in the same transport zone can communicate with each other.
5. Enter the name of an uplink teaming policy.
6. Set **Admin Status** to either **Up** or **Down**.

7 Select a replication mode for the logical switch.

The replication mode (hierarchical two-tier or head) is required for overlay logical switches, but not for VLAN-based logical switches.

Replication Mode	Description
Hierarchical two-tier	The replicator is a host that performs replication of BUM traffic to other hosts within the same VNI. Each host nominates one host tunnel endpoint in every VNI to be the replicator. This is done for each VNI.
Head	Hosts create a copy of each BUM frame and send the copy to each tunnel endpoint it knows about for each VNI.

8 (Optional) Specify a VLAN ID or ranges of VLAN IDs for VLAN tagging.

To support guest VLAN tagging for VMs connected to this switch, you must specify VLAN ID ranges, also called trunk VLAN ID ranges. The logical port will filter packets based on the trunk VLAN ID ranges, and a guest VM can tag its packets with its own VLAN ID based on the trunk VLAN ID ranges.

9 (Optional) Click the **Switching Profiles** tab and select switching profiles.

10 Click **Save**.

In the NSX Manager UI, the new logical switch is a clickable link.

What to do next

Attach VMs to your logical switch. See Connecting a VM to a Logical Switch in Manager Mode.

Connecting a VM to a Logical Switch in Manager Mode

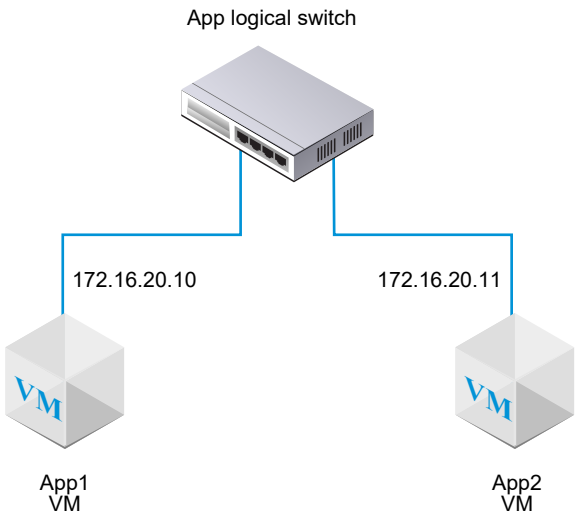
Depending on your host, the configuration for connecting a VM to a logical switch can vary.

The supported hosts that can connect to a logical switch are ESXi hosts in a VMware vCenter and standalone ESXi hosts.

Attach a VM Hosted on vCenter Server to a Logical Switch in Manager Mode

If you have a ESXi host that is managed in vCenter Server, you can access the host VMs through the Web-based vSphere Web Client. In this case, you can use this procedure to attach VMs to NSX logical switches.

The example shown in this procedure shows how to attach a VM called app-vm to a logical switch called app-switch.



The installation-based vSphere Client application does not support attaching a VM to an NSX logical switch. If you do not have the (Web-based) vSphere Web Client, see Attach a VM Hosted on Standalone ESXi to a Logical Switch in Manager Mode.

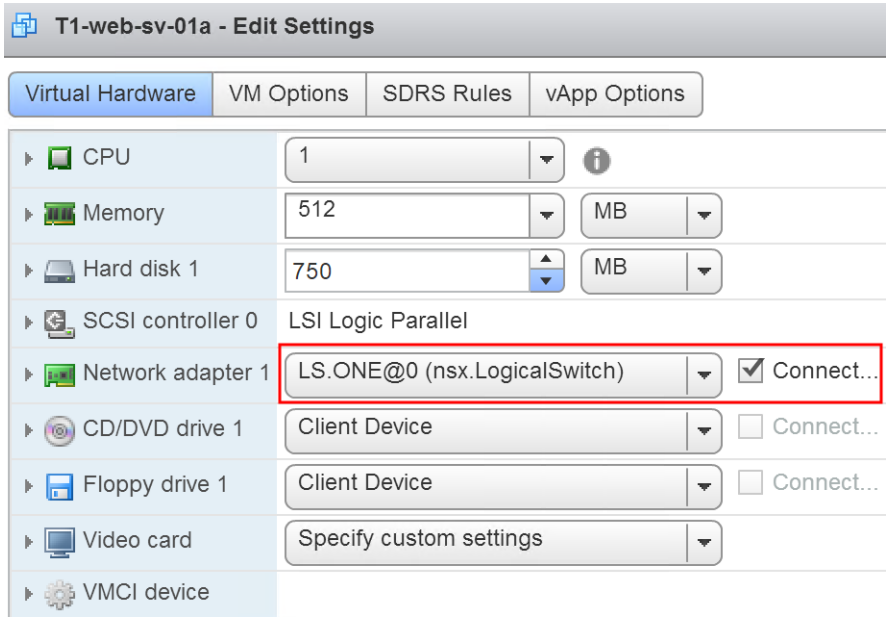
Prerequisites

- The VMs must be hosted on hypervisors that have been added to the NSX fabric.
- The fabric nodes must have NSX management plane (MPA) and NSX control plane (LCP) connectivity.
- The fabric nodes must be added to a transport zone.
- A logical switch must be created.

Procedure

1 In the vSphere Web Client, edit the VM settings, and attach the VM to the NSX logical switch.

For example:



2 Click **OK**.

Results

After attaching a VM to a logical switch, logical switch ports are added to the logical switch. You can view logical switch ports and the VIF attachment ID on the NSX Manager UI. In **Manager mode**, select **Networking > Logical Switches > Ports**.

Use the `GET https://<mgr-ip>/api/v1/logical-ports/` API call to view port details and Admin status for the corresponding VIF attachment ID. To view the Operational status, use the `https://<mgr-ip>/api/v1/logical-ports/<logical-port-id>/status` API call with the appropriate logical port ID.

If two VMs are attached to the same logical switch and have IP addresses configured in the same subnet, they should be able to ping each other.

What to do next

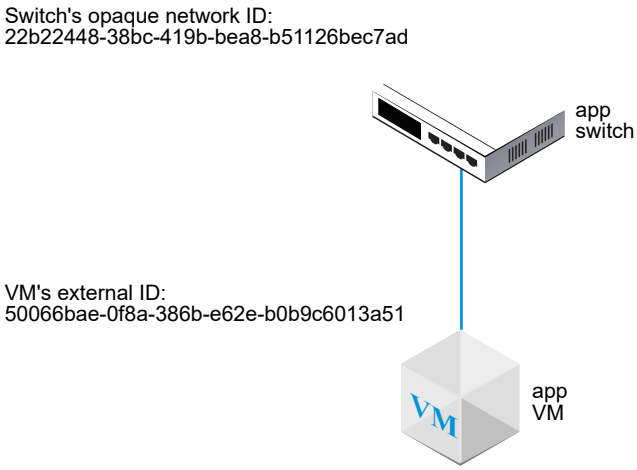
Add a logical router.

You can monitor the activity on the logical switch port to troubleshoot problems. See "Monitor a Logical Switch Port Activity" in the *NSX Administration Guide*.

Attach a VM Hosted on Standalone ESXi to a Logical Switch in Manager Mode

If you have a standalone ESXi host, you cannot access the host VMs through the web-based vSphere Web Client. In this case, you can use this procedure to attach VMs to NSX logical switches.

The example shown in this procedure shows how to attach a VM called app-vm to a logical switch called app-switch.



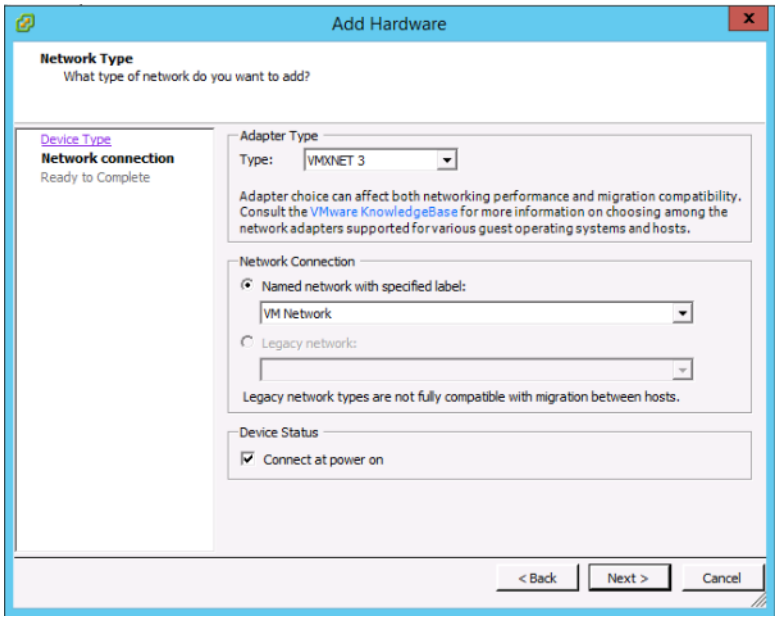
Prerequisites

- The VM must be hosted on hypervisors that have been added to the NSX fabric.
- The fabric nodes must have NSX management plane (MPA) and NSX control plane (LCP) connectivity.
- The fabric nodes must be added to a transport zone.
- A logical switch must be created.
- You must have access to the NSX Manager API.
- You must have write access to the VM's VMX file.

Procedure

1 Using the (install-based) vSphere Client application or some other VM management tool, edit the VM and add a VMXNET 3 Ethernet adapter.

Select any named network. You will change the network connection in a later step.



2 Use the NSX API to issue the `GET https://<nsx-mgr>/api/v1/fabric/virtual-machines/<VM-ID>` API call.

In the results, find the VM's externalId.

For example:

```
GET https://<nsx-mgr>/api/v1/fabric/virtual-machines/60a5a5d5-ea2b-407e-a806-4fdc8468f735

{
  "resource_type": "VirtualMachine",
  "id": "60a5a5d5-ea2b-407e-a806-4fdc8468f735",
  "display_name": "app-vm",
  "compute_ids": [
    "instanceUuid:50066bae-0f8a-386b-e62e-b0b9c6013a51",
    "moIdOnHost:5",
    "externalId:50066bae-0f8a-386b-e62e-b0b9c6013a51",
    "hostLocalId:5",
    "locationId:564dc020-1565-e3f4-f591-ee3953eef3ff",
    "biosUuid:4206f47d-fef7-08c5-5bf7-ea26a4c6b18d"
  ],
  "external_id": "50066bae-0f8a-386b-e62e-b0b9c6013a51",
  "type": "REGULAR",
  "host_id": "cb82b0fa-a8f1-11e5-92a9-6b7d1f8661fa",
```

```
  "local_id_on_host": "5"
}
```

3 Power off and unregister the VM from the host.

You can use your VM management tool or the ESXi CLI, as shown here.

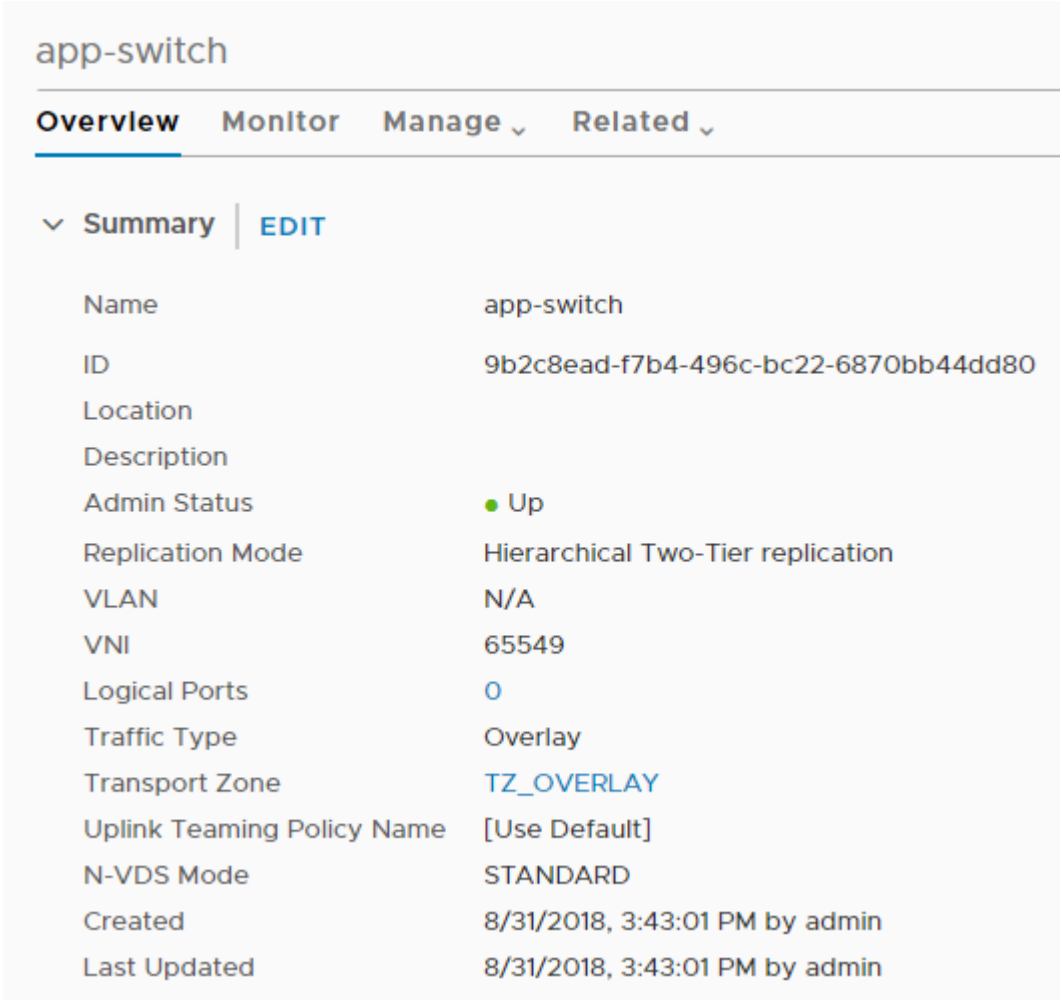
```
[user@host:~] vim-cmd /vmsvc/getallvms
Vmid    Name             File               Guest OS      Version   Annotation
5       app-vm    [ds2] app-vm/app-vm.vmx   ubuntuGuest     vmx-08
8       web-vm    [ds2] web-vm/web-vm.vmx   ubuntu64Guest   vmx-08

[user@host:~] vim-cmd /vmsvc/power.off 5
Powering off VM:

[user@host:~] vim-cmd /vmsvc/unregister 5
```

4 From the NSX Manager UI, get the logical switch ID.

For example:



5 Modify the VM's VMX file.

Delete the **ethernet1.networkName = "<name>"** field and add the following fields:

- ethernet1.opaqueNetwork.id = "<logical switch's ID>"
- ethernet1.opaqueNetwork.type = "nsx.LogicalSwitch"
- ethernet1.externalId = "<VM's externalId>"
- ethernet1.connected = "TRUE"
- ethernet1.startConnected = "TRUE"

For example:

```
OLD
ethernet1.pciSlotNumber = "224"
ethernet1.virtualDev = "vmxnet3"
ethernet1.networkName = "VM Network"
ethernet1.addressType = "vpx"
```

```
ethernet1.generatedAddress = "00:50:56:86:7b:d7"
ethernet1.uptCompatibility = "true"
ethernet1.present = "TRUE"
```

```
NEW
ethernet1.pciSlotNumber = "224"
ethernet1.virtualDev = "vmxnet3"
ethernet1.addressType = "vpx"
ethernet1.generatedAddress = "00:50:56:86:7b:d7"
ethernet1.uptCompatibility = "true"
ethernet1.present = "TRUE"
ethernet1.opaqueNetwork.id = "22b22448-38bc-419b-bea8-b51126bec7ad"
ethernet1.opaqueNetwork.type = "nsx.LogicalSwitch"
ethernet1.externalId = "50066bae-0f8a-386b-e62e-b0b9c6013a51"
ethernet1.connected = "TRUE"
ethernet1.startConnected = "TRUE"
```

6 In the NSX Manager UI, add a logical switch port, and use the VM's externalId for the VIF attachment.

7 Reregister the VM and power it on.

You can use your VM management tool or the ESXi CLI, as shown here.

```
[user@host:~] vim-cmd /solo/register /path/to/file.vmx

For example:
[user@host:~] vim-cmd solo/registervm /vmfs/volumes/355f2049-6c704347/app-vm/app-vm.vmx
9

[user@host:~] vim-cmd /vmsvc/power.on 9
Powering on VM:
```

Results

In the NSX Manager UI in **Manager** mode, select **Networking > Logical Switches > Ports**. Find the VIF attachment ID matching the VM's externalId and make sure that the Admin and Operational status are Up/Up.

If two VMs are attached to the same logical switch and have IP addresses configured in the same subnet, they should be able to ping each other.

What to do next

Add a logical router.

You can monitor the activity on the logical switch port to troubleshoot problems. See "Monitor a Logical Switch Port Activity" in the *NSX Administration Guide*.

Create a Logical Switch Port In Manager Mode

A logical switch has multiple switch ports. A logical switch port connects another network component, a VM, or a container to a logical switch.

If you connect a VM to a logical switch on an ESXi host that is managed by VMware vCenter, a logical switch port is created automatically. For more information about connecting a VM to a logical switch, see Connecting a VM to a Logical Switch in Manager Mode.

For more information about connecting a container to a logical switch, see the *NSX Container Plugin for Kubernetes - Installation and Administration Guide*.

Note The IP address and MAC address bound to a logical switch port for a container are allocated by NSX Manager. Do not change the address binding manually.

To monitor activity on a logical switch port, see Monitor a Logical Switch Port Activity in Manager Mode.

Prerequisites

- Verify that a logical switch is created. See Logical Switches in Manager Mode.
- Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Select **Networking > Logical Switches > Ports > Add**.
3. In the **General** tab, complete the port details.

Option	Description
Name and Description	Enter a name and optionally a description.
Logical Switch	Select a logical switch from the drop-down menu.
Admin Status	Select **Up** or **Down**.
Attachment Type	Select **None** or **VIF**. Select **VIF** if this port is for connecting to a VM.
Attachment ID	If the attachment type is **VIF**, enter the attachment ID.

   Using the API, you can set the attachment type to additional values (`LOGICALROUTER`, `BRIDGEENDPOINT`, `DHCP_SERVICE`, `METADATA_PROXY`, `L2VPN_SESSION`). If the attachment type is DHCP service, metadata proxy, or L2 VPN session, the switching profiles for the port must be the default ones. You cannot use any user-defined profile.

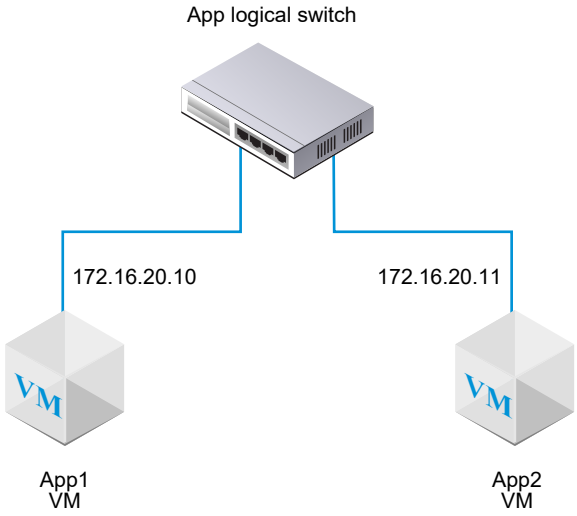
4. (Optional) In the **Switching Profiles** tab, select switching profiles.
5. Click **Save**.

Test Layer 2 Connectivity in Manager Mode

After you successfully set up your logical switch and attach VMs to the logical switch, you can test the network connectivity of the attached VMs.

If your network environment is configured properly, based on the topology the App2 VM can ping the App1 VM.

Figure 25-2. Logical Switch Topology



Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1 Log in to one of the VMs attached to the logical switch using SSH or the VM console.

For example, App2 VM 172.16.20.11.

2 Ping the second VM attached to the logical switch to test connectivity.

```
$ ping -c 2 172.16.20.10
PING 172.16.20.10 (172.16.20.10) 56(84) bytes of data.
64 bytes from 172.16.20.10: icmp_seq=1 ttl=63 time=0.982 ms
64 bytes from 172.16.20.10: icmp_seq=2 ttl=63 time=0.654 ms
64 bytes from 172.16.20.10: icmp_seq=3 ttl=63 time=0.791 ms

--- 172.16.20.10 ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 1990ms
rtt min/avg/max/mdev = 0.654/0.809/0.902/0.104 ms
```

3 (Optional) Identify the problem that causes the ping to fail.

   a Verify that the VM network settings are correct.

   b Verify that the VM network adapter is connected to the correct logical switch.

   c Verify that the logical switch Admin status is UP.

   d From the NSX Manager, select **Networking > Logical Switches > Switches**.

e Click the logical switch and note the UUID and VNI information.

f Run the following commands to troubleshoot the problem.

Command	Description

get logical-switch <vni-or-uuid> arp-table	Displays the ARP table for the specified logical switch.
	Sample output.
	`nsx-manager1> get logical-switch 41866 arp-table`
	`VNI IP MAC Connection-ID`
	`41866 172.16.20.11 00:50:56:b1:70:5e 295422`

get logical-switch <vni-or-uuid> connection-table	Displays the connections for the specified logical switch.
	Sample output.
	`nsx-manager1> get logical-switch 41866 connection-table`
	`Host-IP Port ID`
	`192.168.110.37 36923 295420`
	`192.168.210.53 37883 295421`
	`192.168.210.54 57278 295422`

get logical-switch <vni-or-uuid> mac-table	Displays the MAC table for the specified logical switch.
	Sample output.
	`nsx-manager1> get logical-switch 41866 mac-table`
	`VNI MAC VTEP-IP Connection-ID`
	`41866 00:50:56:86:f2:b2 192.168.250.102 295421`
	`41866 00:50:56:b1:70:5e 192.168.250.101 295422`

get logical-switch <vni-or-uuid> stats	Displays statistics information about the specified logical switch.
	Sample output.
	`nsx-manager1> get logical-switch 41866 stats`
	`update.member 11`
	`update.vtep 11`
	`update.mac 4`
	`update.mac.invalidate 0`
	`update.arp 7`
	`update.arp.duplicate 0`
	`query.mac 2`
	`query.mac.miss 0`
	`query.arp 9`
	`query.arp.miss 6`

get logical-switch <vni-or-uuid> stats-sample	Displays a summary of all logical switch statistics over time.
	Sample output.
	`nsx-manager1> get logical-switch 41866 stats-sample`
	`21:00:00 21:10:00 21:20:00 21:30:00 21:40:00`
	`update.member 0 0 0 0 0`
	`update.vtep 0 0 0 0 0`
	`update.mac 0 0 0 0 0`
	`update.mac.invalidate 0 0 0 0 0`
	`update.arp 0 0 0 0 0`
	`update.arp.duplicate 0 0 0 0 0`

Command	Description
	query.mac 0 0 0 0 0 query.mac.miss 0 0 0 0 0 query.arp 0 0 0 0 0 query.arp.miss 0 0 0 0 0
get logical-switch <vni-or-uuid> vtep	Displays all virtual tunnel end points related to the specified logical switch. Sample output. nsx-manager1> get logical-switch 41866 vtep VNI IP LABEL Segment MAC Connection-ID 41866 192.168.250.102 0x8801 192.168.250.0 00:50:56:65:f5:fc 295421 41866 192.168.250.100 0x1F801 192.168.250.0 02:50:56:00:00:00 295420 41866 192.168.250.101 0x16001 192.168.250.0 00:50:56:64:7c:28 295422

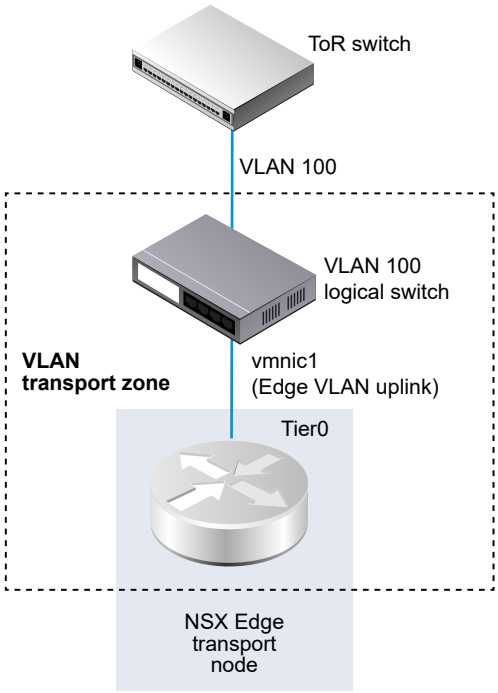
Results

The first VM attached to the logical switch is able to send packets to the second VM.

Create a VLAN Logical Switch for the NSX Edge Uplink in Manager Mode

Edge uplinks go out through VLAN logical switches.

When you are creating a VLAN logical switch, it is important to have in mind a particular topology that you are building. For example, the following simple topology shows a single VLAN logical switch inside of a VLAN transport zone. The VLAN logical switch has VLAN ID 100. This matches the VLAN ID on the TOR port connected to the hypervisor host port used for the Edge's VLAN uplink.



Prerequisites

- To create a VLAN logical switch, you must first create a VLAN transport zone.
- An NSX vSwitch must be added to the NSX Edge. To confirm on an Edge, run the `get host-switches` command. For example:

```
nsx-edge1> get host-switches

Host Switch         : c0a78378-1c20-432a-9e23-ddb34f1c80c9
Switch Name         : hs1
Transport Zone      : c46dcd72-808a-423d-b4cc-8752c33f6b2c
Transport Zone      : 73def985-d122-4b7b-ab6a-a58176dfc32d
Physical Port       : fp-eth0
Uplink Name         : uplink-1
Transport VLAN      : 4096
Default Gateway     : 192.168.150.1
Subnet Mask         : 255.255.255.0
Local VTEP Device   : fp-eth0
Local VTEP IP       : 192.168.150.102
```

- Verify that fabric nodes are successfully connected to the NSX management plane agent (MPA) and the NSX local control plane (LCP).

  In the `GET https://<nsx-mgr>/api/v1/transport-nodes/<transport-node-id>/state` API call, the `state` must be `success`. See the *NSX Installation Guide*.

- Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1 From a browser, log in to an NSX Manager at `https://<nsx-mgr>`.

2 Select **Networking > Logical Switches > Switches > Add**.

3 Type a name for the logical switch.

4 Select a transport zone for the logical switch.

5 Select an uplink teaming policy.

6 For admin status, select **Up** or **Down**.

7 Type a VLAN ID.

   Enter 0 in the VLAN field if there is no VLAN ID for the uplink to the physical TOR.

8 (Optional) Click the **Switching Profiles** tab and select switching profiles.

Results

Note If you have two VLAN logical switches that have the same VLAN ID, they cannot be connected to the same Edge N-VDS (previously known as hostswitch). If you have a VLAN logical switch and an overlay logical switch, and the VLAN ID of the VLAN logical switch is the same as the transport VLAN ID of the overlay logical switch, they also cannot be connected to the same Edge N-VDS.

What to do next

Add a logical router.

Switching Profiles for Logical Switches and Logical Ports

Switching profiles include Layer 2 networking configuration details for logical switches and logical ports. NSX Manager supports several types of switching profiles, and maintains one or more system-defined default switching profiles for each profile type.

The following types of switching profiles are available.

- QoS (Quality of Service)
- Port Mirroring
- IP Discovery
- SpoofGuard
- Switch Security

- MAC Management

Note You cannot edit or delete the default switching profiles in the NSX Manager. You can create custom switching profiles instead.

Before using a default profile, make sure that the settings are what you need them to be. When you create a custom profile, some settings have default values. Do not assume that in the default profile, these settings will have the default values.

Each default or custom switching profile has a unique reserved identifier. You use this identifier to associate the switching profile to a logical switch or a logical port. For example, the default QoS switching profile ID is f313290b-eba8-4262-bd93-fab5026e9495.

A logical switch or logical port can be associated with one switching profile of each type. You cannot have for example, two QoS different switching profiles associated to a logical switch or logical port.

If you do not associate a switching profile type while creating or updating a logical switch, then the NSX Manager associates a corresponding default system-defined switching profile. The children logical ports inherit the default system-defined switching profile from the parent logical switch.

When you create or update a logical switch or logical port you can choose to associate either a default or a custom switching profile. When the switching profile is associated or disassociated from a logical switch the switching profile for the children logical ports is applied based on the following criteria.

- If the parent logical switch has a profile associated with it, the child logical port inherits the switching profile from the parent.
- If the parent logical switch does not have a switching profile associated with it, a default switching profile is assigned to the logical switch and the logical port inherits that default switching profile.
- If you explicitly associate a custom profile with a logical port, then this custom profile overrides the existing switching profile.

Note If you have associated a custom switching profile with a logical switch, but want to retain the default switching profile for one of the child logical port, then you must make a copy of the default switching profile and associate it with the specific logical port.

You cannot delete a custom switching profile if it is associated to a logical switch or a logical port. You can find out whether any logical switches and logical ports are associated with the custom switching profile by going to the Assigned To section of the Summary view and clicking on the listed logical switches and logical ports.

Understanding QoS Switching Profile

QoS provides high-quality and dedicated network performance for preferred traffic that requires high bandwidth. The QoS mechanism does this by prioritizing sufficient bandwidth, controlling

latency and jitter, and reducing data loss for preferred packets even when there is a network congestion. This level of network service is provided by using the existing network resources efficiently.

For this release, shaping and traffic marking namely, CoS and DSCP is supported. The Layer 2 Class of Service (CoS) allows you to specify priority for data packets when traffic is buffered in the logical switch due to congestion. The Layer 3 Differentiated Services Code Point (DSCP) detects packets based on their DSCP values. CoS is always applied to the data packet irrespective of the trusted mode.

NSX trusts the DSCP setting applied by a virtual machine or modifying and setting the DSCP value at the logical switch level. In each case, the DSCP value is propagated to the outer IP header of encapsulated frames. This enables the external physical network to prioritize the traffic based on the DSCP setting on the external header. When DSCP is in the trusted mode, the DSCP value is copied from the inner header. When in the untrusted mode, the DSCP value is not preserved for the inner header.

Note DSCP settings work only on tunneled traffic. These settings do not apply to traffic inside the same hypervisor.

You can use the QoS switching profile to configure the average ingress and egress bandwidth values to set the transmit limit rate. The peak bandwidth rate is used to support burst traffic a logical switch is allowed to prevent congestion on the northbound network links. These settings do not guarantee the bandwidth but help limit the use of network bandwidth. The actual bandwidth you will observe is determined by the link speed of the port or the values in the switching profile, whichever is lower.

The QoS switching profile settings are applied to the logical switch and inherited by the child logical switch port.

Configure a Custom QoS Switching Profile in Manager Mode

You can define the DSCP value and configure the ingress and egress settings to create a custom QoS switching profile.

Prerequisites

- Familiarize yourself with the QoS switching profile concept. See Understanding QoS Switching Profile.
- Identify the network traffic you want to prioritize.
- Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Select **Networking > Logical Switches > Switching Profiles > Add**

3 Select **QoS** and complete the QoS switching profile details.

Option	Description
Name and Description	Assign a name to the custom QoS switching profile. You can optionally describe the setting that you modified in the profile.
Mode	Select either a **Trusted** or **Untrusted** option from the Mode drop-down menu. When you select the Trusted mode the inner header DSCP value is applied to the outer IP header for IP/IPv6 traffic. For non IP/IPv6 traffic, the outer IP header takes the default value. Trusted mode is supported on an overlay-based logical port. The default value is 0. Untrusted mode is supported on overlay-based and VLAN-based logical port. For the overlay-based logical port, the DSCP value of the outbound IP header is set to the configured value irrespective to the inner packet type for the logical port. For the VLAN-based logical port, the DSCP value of IP/IPv6 packet will be set to the configured value. The DSCP values range for untrusted mode is between 0 to 63. **Note** DSCP settings work only on tunneled traffic. These settings do not apply to traffic inside the same hypervisor.
Priority	Set the DSCP value. The priority values range from 0 to 63.
Class of Service	Set the CoS value. CoS is supported on VLAN-based logical port. CoS groups similar types of traffic in the network and each type of traffic is treated as a class with its own level of service priority. The lower priority traffic is slowed down or in some cases dropped to provide better throughput for higher priority traffic. CoS can also be configured for the VLAN ID with zero packet. The CoS values range from 0 to 7, where 0 is the best effort service.
Ingress	Set custom values for the outbound network traffic from the VM to the logical network. You can use the average bandwidth to reduce network congestion. The peak bandwidth rate is used to support burst traffic and the burst size is based on the duration with peak bandwidth. You set burst duration in the burst size setting. You cannot guarantee the bandwidth. However, you can use the Average, Peak, and Burst Size settings to limit network bandwidth. For example, if the average bandwidth is 30 Mbps, peak bandwidth is 60 Mbps, and the allowed duration is 0.1 second, then the burst size is 60 * 1000000 * 0.10/8 = 750000 Bytes. The default value 0 disables rate limiting on the ingress traffic.

Option	Description
Ingress Broadcast	Set custom values for the outbound network traffic from the VM to the logical network based on broadcast. Set custom values for the outbound network traffic from the VM to the logical network based on broadcast. For example, when you set the average bandwidth for a logical switch to 3000 Kbps, peak bandwidth is 6000 Kbps, and the allowed duration is 0.1 second, then the burst size is 6000 * 1000 * 0.10/8 = 75000 Bytes. The default value 0 disables rate limiting on the ingress broadcast traffic.
Egress	Set custom values for the inbound network traffic from the logical network to the VM. The default value 0 disables rate limiting on the egress traffic.

If the ingress, ingress broadcast, and egress options are not configured, the default values are used.

4 Click **Save**.

Results

A custom QoS switching profile appears as a link.

What to do next

Attach this QoS customized switching profile to a logical switch or logical port so that the modified parameters in the switching profile are applied to the network traffic. See Associate a Custom Profile with a Logical Switch in Manager Mode or Associate a Custom Profile with a Logical Port in Manager Mode.

Understanding Port Mirroring Switching Profile

Logical port mirroring lets you replicate and redirect all of the traffic coming in or out of a logical switch port attached to a VM VIF port. The mirrored traffic is sent encapsulated within a Generic Routing Encapsulation (GRE) tunnel to a collector so that all of the original packet information is preserved while traversing the network to a remote destination.

We recommend you use port mirroring only for troubleshooting.

Note Port Mirroring is not recommended for monitoring because when used for longer durations performance is impacted.

Compared to the physical port mirroring, logical port mirroring ensures that all of the VM network traffic is captured. If you implement port mirroring only in the physical network, some of the VM network traffic fails to be mirrored. This happens because communication between VMs residing on the same host never enters the physical network and therefore does not get mirrored. With logical port mirroring you can continue to mirror VM traffic even when that VM is migrated to another host.

The port mirroring process is similar for both VM ports in the NSX domain and ports of physical applications. You can forward the traffic captured by a workload connected to a logical network and mirror that traffic to a collector. The IP address should be reachable from the guest IP address on which the VM is hosted. This process is also true for physical applications connected to gateway nodes.

Configure a Custom Port Mirroring Switching Profile in Manager Mode

You can create a custom port mirroring switching profile with a different destination and key value.

Prerequisites

- Familiarize yourself with the port mirroring switching profile concept. See Understanding Port Mirroring Switching Profile.
- Identify the IP address of the destination logical port ID you want to redirect network traffic to.
- Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Select **Networking > Logical Switches > Switching Profiles > Add**
3. Select **Port Mirroring** and complete the port mirroring switching profile details.

Option	Description
Name and Description	Assign a name to the custom port mirroring switching profile. You can optionally describe the setting you modified to customize this profile.
Direction	Select an option from the drop-down menu to use this source for **Ingress**, **Egress**, or **Bidirectional** traffic. Ingress is the outbound network traffic from the VM to the logical network. Egress is the inbound network traffic from the logical network to the VM. Bidirectional is the two-way of traffic from the VM to the logical network and from the logical network to the VM. This is the default option.
Packet Truncation	Optional. The range is 60 - 65535.

Option	Description
Key	Enter a random 32-bit value to identify mirrored packets from the logical port. This Key value is copied to the Key field in the GRE header of each mirror packet. If the Key value is set to 0, the default definition is copied to the Key field in the GRE header. The default 32-bit value is made of the following values. ■ The first 24-bit is a VNI value. VNI is part of the IP header of encapsulated frames. ■ The 25th bit indicates if the first 24-bit is a valid VNI value. One represents a valid value and zero represents an invalid value. ■ The 26th bit indicates the direction of the mirrored traffic. One represents an ingress direction and zero represents an egress direction. ■ The remaining six bits are not used.
Destinations	Enter the destination ID of the collector for the mirroring session. The destination IP address ID can only be an IPv4 address within the network or a remote IPv4 address not managed by NSX. You can add up to three destination IP addresses separated by a comma.

4 Click **Save**.

Results

A custom port mirroring switching profile appears as a link.

What to do next

Attach the switching profile to a logical switch or logical port. See Associate a Custom Profile with a Logical Switch in Manager Mode or Associate a Custom Profile with a Logical Port in Manager Mode.

Verify that the customized port mirroring switching profile works. See Verify Custom Port Mirroring Switching Profile.

Verify Custom Port Mirroring Switching Profile

Before you start using the custom port mirroring switching profile, verify that the customization works properly.

Prerequisites

- Verify that the custom port mirroring switching profile is configured. See Configure a Custom Port Mirroring Switching Profile in Manager Mode.
- Verify that the customized port mirroring switching profile is attached to a logical switch. See Associate a Custom Profile with a Logical Switch in Manager Mode.

Procedure

1 Locate two VMs with VIF attachments to the logical port configured for port mirroring.

   For example, VM1 10.70.1.1 and VM2 10.70.1.2 have VIF attachments and they are located in the same logical network.

2 Run the `tcpdump` command on a destination IP address.

   `sudo tcpdump -n -i eth0 dst host destination_IP_addres and proto gre`

   For example, the destination IP address is 10.24.123.196.

3 Log in to the first VM and ping the second VM to verify that the corresponding ECHO requests and replies are received at the destination address.

What to do next

Attach this port mirroring customized switching profile to a logical switch so that the modified parameters in the switching profile are applied to the network traffic. See Associate a Custom Profile with a Logical Switch in Manager Mode.

Understanding IP Discovery Switching Profile

IP Discovery uses DHCP and DHCPv6 snooping, ARP (Address Resolution Protocol) snooping, ND (Neighbor Discovery) snooping, and VM Tools to learn MAC and IP addresses.

The discovered MAC and IP addresses are used to achieve ARP/ND suppression, which minimizes traffic between VMs connected to the same logical switch. The addresses are also used by the SpoofGuard and distributed firewall (DFW) components. DFW uses the address bindings to determine the IP address of objects in firewall rules.

DHCP/DHCPv6 snooping inspects the DHCP/DHCPv6 packets exchanged between the DHCP/DHCPv6 client and server to learn the IP and MAC addresses.

ARP snooping inspects the outgoing ARP and GARP (gratuitous ARP) packets of a VM to learn the IP and MAC addresses.

VM Tools is software that runs on an ESXi-hosted VM and can provide the VM's configuration information including MAC and IP or IPv6 addresses. This IP discovery method is available for VMs running on ESXi hosts only.

ND snooping is the IPv6 equivalent of ARP snooping. It inspects neighbor solicitation (NS) and neighbor advertisement (NA) messages to learn the IP and MAC addresses.

Duplicate address detection checks whether a newly discovered IP address is already present on the realized binding list for a different port. This check is performed for ports on the same segment. If a duplicate address is detected, the newly discovered address is added to the discovered list, but is not added to the realized binding list. All duplicate IPs have an associated discovery timestamp. If the IP that is on the realized binding list is removed, either by adding it to the ignore binding list or by disabling snooping, the duplicate IP with the oldest timestamp is moved to the realized binding list. The duplicate address information is available through an API call.

By default, the discovery methods ARP snooping and ND snooping operate in a mode called trust on first use (TOFU). In TOFU mode, when an address is discovered and added to the realized bindings list, that binding remains in the realized list forever. TOFU applies to the first 'n' unique <IP, MAC, VLAN> bindings discovered using ARP/ND snooping, where 'n' is the binding limit that you can configure. You can disable TOFU for ARP/ND snooping. The methods will then operate in trust on every use (TOEU) mode. In TOEU mode, when an address is discovered, it is added to the realized bindings list and when it is deleted or expired, it is removed from the realized bindings list. DHCP snooping and VM Tools always operate in TOEU mode

For each port, NSX Manager maintains an ignore bindings list, which contains IP addresses that cannot be bound to the port. If you navigate to **Networking > Logical Switches > Ports** in **Manager** mode, and select a port, you can add discovered bindings to the ignore bindings list. You can also delete an existing discovered or realized binding by copying it to **Ignore Bindings**.

Note TOFU is not the same as SpoofGuard, and it does not block traffic in the same way as SpoofGuard. For more information, see Understanding SpoofGuard Segment Profile.

For Linux VMs, the ARP flux problem might cause ARP snooping to obtain incorrect information. The problem can be prevented with an ARP filter. For more information, see http://linux-ip.net/html/ether-arp.html#ether-arp-flux.

Configure IP Discovery Switching Profile in Manager Mode

NSX has several default IP Discovery switching profiles. You can also create additional ones.

Prerequisites

- Familiarize yourself with the IP Discovery switching profile concepts. See Understanding IP Discovery Switching Profile.
- Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Select **Networking > Logical Switches > Switching Profiles > Add**.
3. Select **IP Discovering** and specify the IP Discovery switching profile details.

Option	Description
Name and Description	Enter a name and optionally a description.
ARP Snooping	For an IPv4 environment. Applicable if VMs have static IP addresses.
ARP Binding Limit	The maximum number of IPv4 IP addresses that can be bound to a port. The minimum value allowed is 1 (the default) and the maximum is 256.
ARP ND Binding Limit Timeout	The timeout value, in minutes, for IP addresses in the ARP/ND binding table if TOFU is disabled. If an address times out, a newly discovered address replaces it.

Option	Description
DHCP Snooping	For an IPv4 environment. Applicable if VMs have IPv4 addresses.
DHCP V6 Snooping	For an IPv6 environment. Applicable if VMs have IPv6 addresses.
VM Tools	Available for ESXi-hosted VMs only.
VM Tools for IPv6	Available for ESXi-hosted VMs only.
Neighbor Discovery Snooping	For an IPv6 environment. Applicable if VMs have static IP addresses.
Neighbor Discovery Binding Limit	The maximum number of IPv6 addresses that can be bound to a port.
Trust on First Use	Applicable to ARP and ND snooping.
Duplicate IP Detection	For all snooping methods and both IPv4 and IPv6 environments.

4 Click **Add**.

What to do next

Attach this IP Discovery customized switching profile to a logical switch or logical port so that the modified parameters in the switching profile are applied to the network traffic. See Associate a Custom Profile with a Logical Switch in Manager Mode or Associate a Custom Profile with a Logical Port in Manager Mode.

Configure IP Discovery Segment Profile on Groups

Configuring IP Discovery segment profiles on a group allows a Security Administrator to configure IP discovery profile parameters, and apply them to group members.

Configuring The following static and dynamic group members are supported:

- Segment
- Segment Port
- VM
- Groups
- Mix of the above

Profiles on groups only apply if the default profile is applied to the segment or segment port:

Custom Group Profile	Custom Profile on Segment (S) and Segment Port(SP)	Effective Profile on Port
Custom	Default (S), Default (SP)	Custom
Custom 1	Default (S), Custom 2 (SP)	Custom 2
Custom 1	Custom 2 (S), Default (SP)	Custom 2
Custom 1	Custom 2 (S), Custom 3 (SP)	Custom 3

Each time a profile is applied to a group a sequence number is specified. If a member is present in multiple groups, the group with the lower sequence number has higher priority.

Discovery Profile Binding Map API

Method	API	Resource Type
PUT, PATCH, GET, DELETE	`/infra/domains/<domain-id>/groups/<group-id>/discovery-profile-binding-maps/<binding-map-id>`	DiscoveryProfileBindingMap
GET	`/infra/domains/<domain-id>/groups/<group-id>/discovery-profile-binding-maps`	DiscoveryProfileBindingMapListResult

Parameters for DiscoveryProfileBindingMap

Field	Type	Description
profile_path	Policy Path	Required
sequence_number	Integer	Required. Sequence number is used to resolve conflicts when two profiles are applied to the same segment or segment port. The low sequence number has higher precedence.

API for Segments and Ports

Method	API	Resource Type

GET	`/infra/tier-1s/<tier-1-id>/segments/<segment-id>/effective-profiles`
	`/infra/segments/<segment-id>/effective-profiles`
	`/infra/segments/<segment-id>/effective-profiles`
	`/infra/segments/<segment-id>/ports/<port-id>`
		EffectiveProfilesResponse

Example Request

POST https://{{policy-ip}}/policy/api/v1/infra/domains/default/groups/TestGroup/discovery-profile-binding-maps/ipdmap

```
{
    "profile_path" : "/infra/ip-discovery-profiles/ip-discovery-custom-profile-1",
    "sequence_number" :"10"
}
```

Understanding SpoofGuard

SpoofGuard helps prevent a form of malicious attack called "web spoofing" or "phishing." A SpoofGuard policy blocks traffic determined to be spoofed.

SpoofGuard is a tool that is designed to prevent virtual machines in your environment from altering their existing IP address. In the instance that a virtual machine's IP address does not match the IP address on the corresponding logical port and switch address binding in SpoofGuard, the virtual machine's vNIC is prevented from accessing the network entirely. SpoofGuard can be configured at the port or switch level. There are several reasons SpoofGuard might be used in your environment:

- Preventing a rogue virtual machine from assuming the IP address of an existing VM.
- Ensuring the IP addresses of virtual machines cannot be altered without intervention – in some environments, it's preferable that virtual machines cannot alter their IP addresses without proper change control review. SpoofGuard facilitates this by ensuring that the virtual machine owner cannot simply alter the IP address and continue working unimpeded.
- Guaranteeing that distributed firewall (DFW) rules will not be inadvertently (or deliberately) bypassed – for DFW rules created utilizing IP sets as sources or destinations, the possibility always exists that a virtual machine could have it's IP address forged in the packet header, thereby bypassing the rules in question.

NSX SpoofGuard configuration covers the following:

- MAC SpoofGuard - authenticates MAC address of packet
- IP SpoofGuard - authenticates MAC and IP addresses of packet
- Dynamic Address Resolution Protocol (ARP) inspection, that is, ARP and Gratuitous Address Resolution Protocol (GARP) SpoofGuard and Neighbor Discovery (ND) SpoofGuard validation are all against the MAC source, IP Source and IP-MAC source mapping in the ARP/GARP/ND payload.

At the port level, the allowed MAC/VLAN/IP allow-list is provided through the Address Bindings property of the port. When the virtual machine sends traffic, it is dropped if its IP/MAC/VLAN does not match the IP/MAC/VLAN properties of the port. The port level SpoofGuard deals with traffic authentication, i.e. is the traffic consistent with VIF configuration.

At the switch level, the allowed MAC/VLAN/IP allow-list is provided through the Address Bindings property of the switch. This is typically an allowed IP range/subnet for the switch and the switch level SpoofGuard deals with traffic authorization.

Traffic must be permitted by port level AND switch level SpoofGuard before it will be allowed into switch. Activating or deactivating port and switch level SpoofGuard, can be controlled using the SpoofGuard switch profile.

Configure Port Address Bindings in Manager Mode

Address bindings specify the IP and MAC address of a logical port and are used to specify the port allow-list in SpoofGuard.

With port address bindings you'll specify the IP and MAC address, and VLAN if applicable, of the logical port. When SpoofGuard is enabled, it ensures that the specified address bindings are enforced in the data path. In addition to SpoofGuard, port address bindings are used for DFW rule translations.

Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1. In NSX Manager, select to **Networking** > **Logical Switches** > **Ports**.
2. Click the logical port to which you want apply address binding.

   The logical port summary appears.
3. In the **Overview** tab, expand **Address Bindings** > **Manual Bindings** .
4. Click **Add**.

   The Add Address Binding dialogue box appears.
5. Specify the IP (IPv4 address, IPv4 subnet, IPv6 address, or IPv6 subnet) and MAC address of the logical port to which you want to apply address binding. For example, for IPv6, 2001::/64 is an IPv6 subnet, 2001::1 is a host IP, whereas 2001::1/64 is an invalid input. You can also specify a VLAN ID.
6. Click **Add**.

What to do next

Use the port address bindings when you Configure a SpoofGuard Switching Profile in Manager Mode.

Configure a SpoofGuard Switching Profile in Manager Mode

When SpoofGuard is configured, if the IP address of a virtual machine changes, traffic from the virtual machine may be blocked until the corresponding configured port/switch address bindings are updated with the new IP address.

Enable SpoofGuard for the port group(s) containing the guests. When enabled for each network adapter, SpoofGuard inspects packets for the prescribed MAC and its corresponding IP address.

Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Select **Networking** > **Logical Switches** > **Switching Profiles** > **Add**.
3. Select **Spoof Guard**.
4. Enter a name and optionally a description.

5 To enable port level SpoofGuard, set **Port Bindings** to **Enabled**.

6 Click **Add**.

Results

A new switching profile has been created with a SpoofGuard Profile.

What to do next

Associate the SpoofGuard profile with a logical switch or logical port. See Associate a Custom Profile with a Logical Switch in Manager Mode or Associate a Custom Profile with a Logical Port in Manager Mode.

Understanding Switch Security Switching Profile

Switch security provides stateless Layer2 and Layer 3 security by checking the ingress traffic to the logical switch and dropping unauthorized packets sent from VMs by matching the IP address, MAC address, and protocols to a set of allowed addresses and protocols. You can use switch security to protect the logical switch integrity by filtering out malicious attacks from the VMs in the network.

You can configure the Bridge Protocol Data Unit (BPDU) filter, DHCP Snooping, DHCP server block, and rate limiting options to customize the switch security switching profile on a logical switch.

Configure a Custom Switch Security Switching Profile in Manager Mode

You can create a custom switch security switching profile with MAC destination addresses from the allowed BPDU list and configure rate limiting.

Prerequisites

- Familiarize yourself with the switch security switching profile concept. See Understanding Switch Security Switching Profile.
- Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Networking** > **Logical Switches**.

3 Click the **Switching Profiles** tab.

4 Click **Add** and select **Switch Security**.

5 Complete the switch security profile details.

Option	Description
Name and Description	Assign a name to the custom switch security profile. You can optionally describe the setting that you modified in the profile.
BPDU Filter	Toggle the **BPDU Filter** button to enable BPDU filtering. Disabled by default. When the BPDU filter is enabled, all of the traffic to BPDU destination MAC address is blocked. The BPDU filter when enabled also disables STP on the logical switch ports because these ports are not expected to take part in STP.
BPDU Filter Allow List	Click the destination MAC address from the BPDU destination MAC addresses list to allow traffic to the permitted destination. You must enable **BPDU Filter** to be able to select from this list.
DHCP Filter	Toggle the **Server Block** button and **Client Block** button to enable DHCP filtering. Both are disabled by default. DHCP Server Block blocks traffic from a DHCP server to a DHCP client. Note that it does not block traffic from a DHCP server to a DHCP relay agent. DHCP Client Block prevents a VM from acquiring a DHCP IP address by blocking DHCP requests.
DHCPv6 Filter	Toggle the **V6 Server Block** button and **V6 Client Block** button to enable DHCP filtering. Both are disabled by default. DHCPv6 Server Block blocks traffic from a DHCPv6 server to a DHCPv6 client. Note that it does not block traffic from a DHCP server to a DHCP relay agent. Packets whose UDP source port number is 547 are filtered. DHCPv6 Client Block prevents a VM from acquiring a DHCP IP address by blocking DHCP requests. Packets whose UDP source port number is 546 are filtered.
Block Non-IP Traffic	Toggle the **Block Non-IP Traffic** button to allow only IPv4, IPv6, ARP, and BPDU traffic. The rest of the non-IP traffic is blocked. The permitted IPv4, IPv6, ARP, GARP and BPDU traffic is based on other policies set in address binding and SpoofGuard configuration. By default, this option is disabled to allow non-IP traffic to be handled as regular traffic.
RA Guard	Toggle the **RA Guard** button to filter out ingress IPv6 router advertisements. ICMPv6 type 134 packets are filtered out. This option is enabled by default.
Rate Limits	Set a rate limit for broadcast and multicast traffic. This option is enabled by default. Rate limits can be used to protect the logical switch or VMs from events such as broadcast storms. To avoid any connectivity problems, the minimum rate limit value must be >= 10 pps.

6 Click **Add**.

Results

A custom switch security profile appears as a link.

What to do next

Attach this switch security customized switching profile to a logical switch or logical port so that the modified parameters in the switching profile are applied to the network traffic. See Associate a Custom Profile with a Logical Switch in Manager Mode or Associate a Custom Profile with a Logical Port in Manager Mode.

Understanding MAC Management Switching Profile

The MAC management switching profile supports two functionalities: MAC learning and MAC address change.

The MAC address change feature allows a VM to change its MAC address. A VM connected to a port can run an administrative command to change the MAC address of its vNIC and still send and receive traffic on that vNIC. In the default MAC management switching profile, this property is enabled by default.

MAC learning provides network connectivity to deployments where multiple MAC addresses are configured behind one vNIC, for example, in a nested hypervisor deployment where an ESXi VM runs on an ESXi host and multiple VMs run inside the ESXi VM. Without MAC learning, when the ESXi VM's vNIC connects to a switch port, its MAC address is static. VMs running inside the ESXi VM do not have network connectivity because their packets have different source MAC addresses. With MAC learning, the vSwitch inspects the source MAC address of every packet coming from the vNIC, learns the MAC address and allows the packet to go through. If a MAC address that is learned is not used for a certain period of time, it is removed. This aging property is not configurable.

MAC learning also supports unknown unicast flooding. Normally, when a packet that is received by a port has an unknown destination MAC address, the packet is dropped. With unknown unicast flooding enabled, the port floods unknown unicast traffic to every port on the switch that has MAC learning and unknown unicast flooding enabled. This property is enabled by default, but only if MAC learning is enabled.

The number of MAC addresses that can be learned is configurable. The maximum value is 4096, which is the default. You can also set the policy for when the limit is reached. The options are:

- **Drop** - Packets from an unknown source MAC address are dropped. Packets inbound to this MAC address will be treated as unknown unicast. The port will receive the packets only if it has unknown unicast flooding enabled.
- **Allow** - Packets from an unknown source MAC address are forwarded although the address will not be learned. Packets inbound to this MAC address will be treated as unknown unicast. The port will receive the packets only if it has unknown unicast flooding enabled.

If you enable MAC learning or MAC address change, to improve security, configure SpoofGuard as well.

Configure MAC Management Switching Profile in Manager Mode

You can create a MAC management switching profile to manage MAC addresses.

Prerequisites

- Familiarize yourself with the MAC management switching profile concept. See Understanding MAC Management Switching Profile.
- Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Select **Networking > Logical Switches > Switching Profiles > Add**.
3. Select **MAC Management** and complete the MAC management profile details.

Option	Description
Name and Description	Assign a name to the MAC management profile. You can optionally describe the setting that you modified in the profile.
MAC Change	Enable or disable the MAC address change feature. The default is disabled.
Status	Enable or disable the MAC learning feature. The default is disabled.
Unknown Unicast Flooding	Enable or disable the unknown unicast flooding feature. The default is enabled. This option is available if you enable MAC learning
MAC Limit	Set the maximum number of MAC addresses. The default is 4096. This option is available if you enable MAC learning
MAC Limit Policy	Select **Allow** or **Drop**. The default is **Allow**. This option is available if you enable MAC learning

4. Click **Add**.

What to do next

Attach the switching profile to a logical switch or logical port. See Associate a Custom Profile with a Logical Switch in Manager Mode or Associate a Custom Profile with a Logical Port in Manager Mode.

Associate a Custom Profile with a Logical Switch in Manager Mode

You can associate a custom switching profile to a logical switch so that the profile applies to all the ports on the switch.

When custom switching profiles are attached to a logical switch they override existing default switching profiles. The custom switching profile is inherited by children logical switch ports.

Note If you have associated a custom switching profile with a logical switch, but want to retain the default switching profile for one of the child logical switch port, then you must make a copy of the default switching profile and associate it with the specific logical switch port.

Prerequisites

- Verify that a logical switch is configured. See Create a Logical Switch in Manager Mode.
- Verify that a custom switching profile is configured. See Switching Profiles for Logical Switches and Logical Ports.
- Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Select **Networking > Logical Switches > Switches**.
3. Click the logical switch to apply the custom switching profile.
4. Click the **Manage** tab.
5. Select the custom switching profile type from the drop-down menu.
   - **QoS**
   - **Port Mirroring**
   - **IP Discovering**
   - **SpoofGuard**
   - **Switch Security**
   - **MAC Management**
6. Click **Change.**
7. Select the previously created custom switching profile from the drop-down menu.
8. Click **Save**.

   The logical switch is now associated with the custom switching profile.
9. Verify that the new custom switching profile with the modified configuration appears under the **Manage** tab.
10. (Optional) Click the **Related** tab and select **Ports** from the drop-down menu to verify that the custom switching profile is applied to child logical ports.

What to do next

If you do not want to use the inherited switching profile from a logical switch, you can apply a custom switching profile to the child logical switch port. See Associate a Custom Profile with a Logical Port in Manager Mode.

Associate a Custom Profile with a Logical Port in Manager Mode

A logical port provides a logical connection point for a VIF, a patch connection to a router, or a Layer 2 gateway connection to an external network. Logical ports also expose switching profiles, port statistics counters, and a logical link status.

You can change the inherited switching profile from the logical switch to a different custom switching profile for the child logical port.

Prerequisites

- Verify that a logical port is configured. See Connecting a VM to a Logical Switch in Manager Mode.
- Verify that a custom switching profile is configured. See Switching Profiles for Logical Switches and Logical Ports.
- Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Select **Networking > Logical Switches > Ports**.
3. Click the logical port to apply the custom switching profile.
4. Click the **Manage** tab.
5. Select the custom switching profile type from the drop-down menu.
   - **QoS**
   - **Port Mirroring**
   - **IP Discovering**
   - **SpoofGuard**
   - **Switch Security**
   - **MAC Management**
6. Click **Change**.
7. Select the previously created custom switching profile from the drop-down menu.
8. Click **Save**.

   The logical port is now associated with the custom switching profile.
9. Verify that the new custom switching profile with the modified configuration appears under the **Manage** tab.

What to do next

You can monitor the activity on the logical switch port to troubleshoot problems. See Monitor a Logical Switch Port Activity in Manager Mode.

Edge Bridging in Manager Mode: Extending Overlay Segments to VLAN

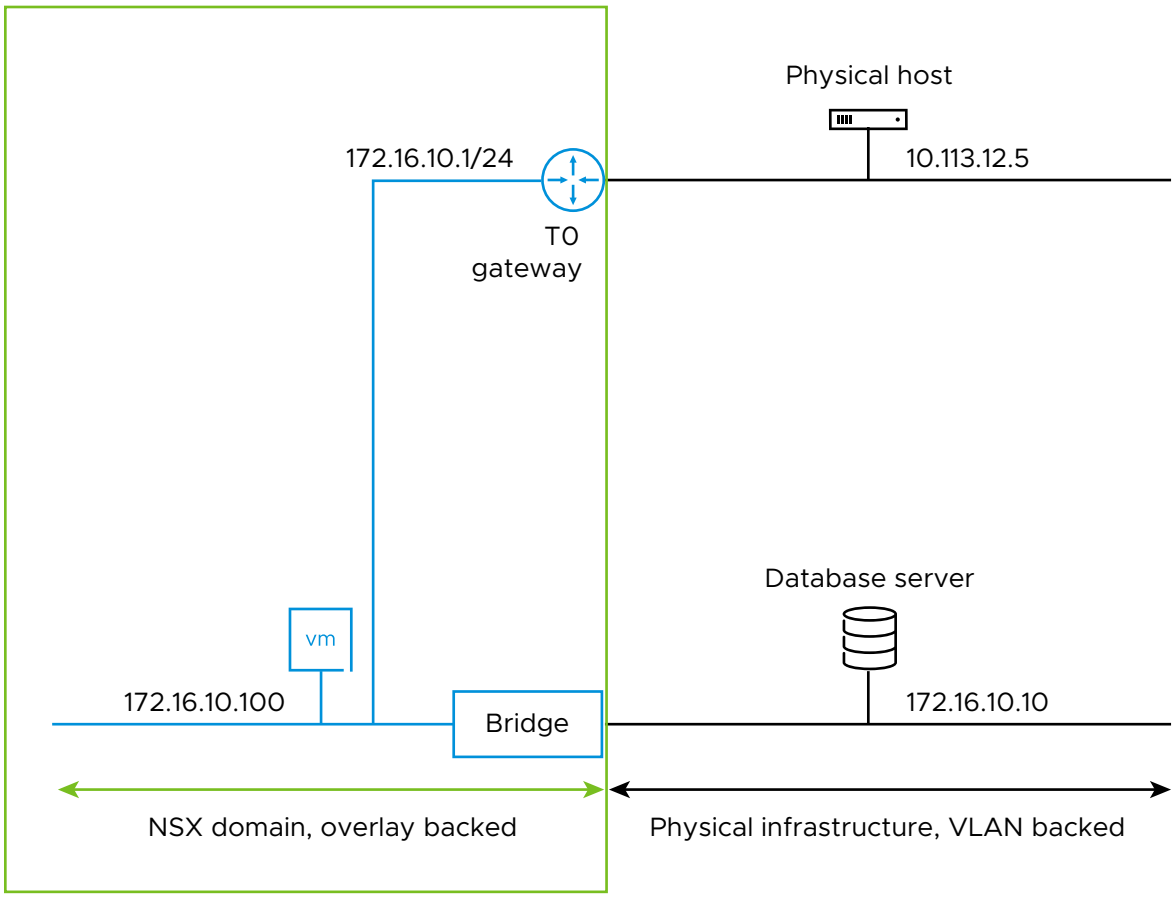
Workloads attached to overlay segments typically communicate at layer 3 with physical devices outside of the NSX-T Data Center domain, through tier-0 gateways instantiated on NSX Edge. However, there are some scenarios where layer 2 connectivity is required between virtual machines in NSX-T Data Center and physical devices.

Some examples are:

- Migration from physical to virtual, or virtual to virtual.
- Integration of a physical appliance that provides services to a segment, like an external load balancer.
- Connection to a database server that requires layer 2 adjacency to its virtual machine clients.

For that purpose, on the top of the gateway service, NSX Edge can also run a bridge service. The following diagram represents those two options: the virtual machine in the bottom left corner has layer 3 connectivity through a gateway to the physical host, and layer 2 connectivity through a bridge to the database server. It is possible to both route and bridge a segment. In fact, it is possible to use the tier 0 gateway in this diagram as a default gateway for the database server.

Figure 25-3. NSX VM Bridge and Gateway Communication



The NSX Edge bridge, like the gateway, is supported for long term deployments, even if it is often used as a temporary solution during migrations.

The bridge functionality extends an overlay segment into a VLAN, identified by a VLAN ID on an uplink of the NSX Edge where the bridge is running. Typically, two redundant active and standby bridges get deployed on separate edges as part of the same edge cluster. There is no active/active redundancy possible. Setting up the bridge functionality involves the following configuration steps:

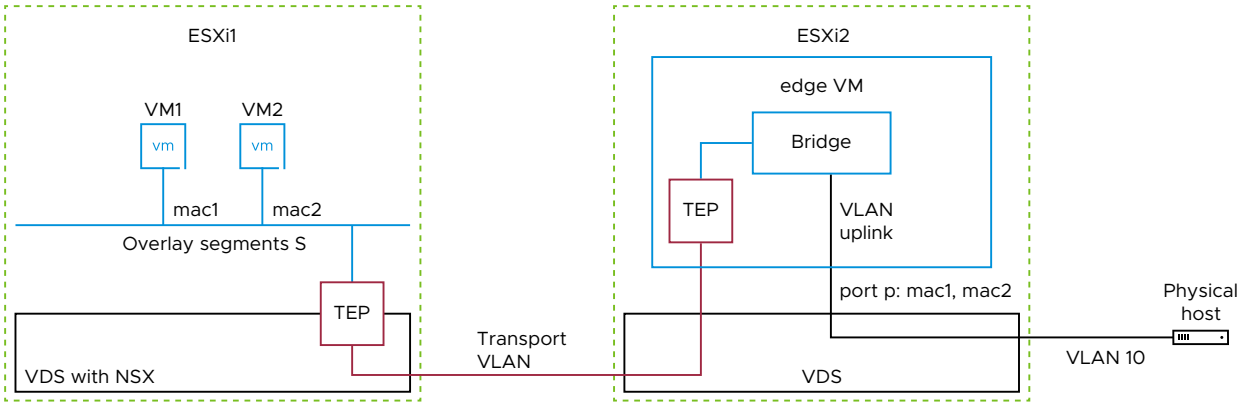
- Make sure that the NSX Edge is suitable for hosting the bridge service. The bridge adds a few constraints to the deployment of an edge in a VM form factor.
- Identify the NSX Edges that run the bridge service. A bridge profile statically designates the edge responsible for running the active bridge and optionally designates a second edge hosting the standby bridge.
- Lastly associate an overlay segment to a VLAN ID or IDs and a bridge profile. This results in the creation of the appropriate active/standby bridges on the edges specified in the bridge profile, that extend at layer 2 the overlay segment to the VLAN or VLANs identified by the VLAN IDs.

Configure an Edge VM for Bridging in Manager Mode

There are no specific constraints to configure bridging on a bare metal edge. However, if you are planning to run a bridge on an NSX Edge VM, use this section to understand the specific configuration to perform in the vSphere infrastructure.

As an example, our scenario includes two virtual machines, VM1 and VM2, on transport node ESXi 1 attached to an overlay segment S. The VMs can communicate at layer 2 with the physical host on the right side of the diagram thanks to a bridge instantiated on the edge VM running on ESXi host 2. The TEP (tunnel end point) on ESXi 1 encapsulates the traffic from VM1/VM2 and forwards it to the TEP of the edge VM. Then the bridge unencapsulates the traffic and sends it tagged with VLAN ID 10 on its VLAN uplink. Then the traffic gets switched to the physical host.

Figure 25-4. Edge VM Bridging



Option 1: Edge VM is on a VSS portgroup

This option is for when the Edge VM is connected to a VSS (vSphere Standard Switch). You must enable promiscuous mode and forged transmit.

- Set promiscuous mode on the portgroup.
- Allow forged transmit on the portgroup.
- Run the following command to enable reverse filter on the ESXi host where the Edge VM is running:

  ```
  esxcli system settings advanced set -o /Net/ReversePathFwdCheckPromisc -i 1
  ```

  Then disable and enable promiscuous mode on the portgroup with the following steps:

  - Edit the portgroup's settings.
  - Disable promiscuous mode and save the settings.
  - Edit the portgroup's settings again.
  - Enable promiscuous mode and save the settings.
- Do not have other port groups in promiscuous mode on the same host sharing the same set of VLANs.

- Avoid running other VMs attached to the portgroup in promiscuous mode on the same host, as the traffic gets replicated to all those VMs and affect performance.

Option 2a: Edge VM is on a VDS 6.6.0 (or later) portgroup

This option is for when the Edge VM is connected to a VDS (vSphere Distributed Switch). You must be running ESXi 6.7 or later, and VDS 6.6.0 or later.

- Enable MAC learning with the option “allow unicast flooding” on the distributed portgroup.

  Starting with vSphere 8.0, you can enable the Mac Learning UI option in the distributed portgroup configuration. For previous releases, you need to use the VIM API `DVSMacLearningPolicy` and setting `allowUnicastFlooding` to `true`.

Option 2b: Edge VM is on a VDS 6.5.0 (or later) portgroup

This option is for when the Edge VM is connected to a VDS. You enable promiscuous mode and forged transmit.

- Set promiscuous mode on the distributed portgroup.
- Allow forged transmit on the distributed portgroup.
- Run the following command to enable reverse filter on the ESXi host where the Edge VM is running:

  ```
  esxcli system settings advanced set -o /Net/ReversePathFwdCheckPromisc -i 1
  ```

  Then disable and enable promiscuous mode on the distributed portgroup with the following steps:

  - Edit the distributed portgroup's settings.
  - Disable promiscuous mode and save the settings.
  - Edit the distributed portgroup's settings again.
  - Enable promiscuous mode and save the settings.
- Do not have other distributed port groups in promiscuous mode on the same host sharing the same set of VLANs.
- Avoid running other VMs attached to the distributed portgroup in promiscuous mode on the same host, as the traffic gets replicated to all those VMs and affects performance.

Option 3: Edge VM is connected to an NSX segment

If the Edge is deployed on a host with NSX installed, it can connect to a VLAN segment and use MAC Learning, which is the preferred configuration option.

- Create a new MAC Discovery segment profile by navigating to **Networking > Segments > Profiles**.
  - Click **Add Segment Profile > MAC Discovery**.

- Enable **MAC Learning**. This will also enable **Unknown Unicast Flooding**. Keep the flooding option enabled for bridging to work in all scenarios.
- Click **Save**.

- Edit the segment used by the Edge by navigating to **Networking > Segments**.
  - Click the menu icon (3 dots) and select **Edit**.
  - Expand the **Segment Profiles** section, then set the **MAC Discovery** profile to the one created above.

Note If you bridge a segment to VLAN 0 and you use a distributed router on this segment, the gateway might not route VLAN 0 traffic when using MAC learning. In this scenario, avoid option 3. Avoid option 2a if the edge VM is attached to the distributed portgroup of a VDS prepared for NSX for vSphere.

Create an Edge Bridge Profile in Manager Mode

The edge bridge profile is a template for instantiating bridges. In the template, you define a primary edge, an optional backup edge from the same edge cluster as the primary, and a failover mode, preemptive or non-preemptive.

Select the Edges

The preference is to use the primary edge for running the active bridge, the bridge forwarding traffic between overlay segment and VLAN. The standby bridge, that typically runs on the backup edge, does not forward any traffic.

You can instantiate multiple bridges from the same bridge profile. As a result, in most cases, few edge bridge profiles are required. For example, if you plan to use two edges (edge1 and edge2) for bridging, you might want to create two edge bridge profiles:

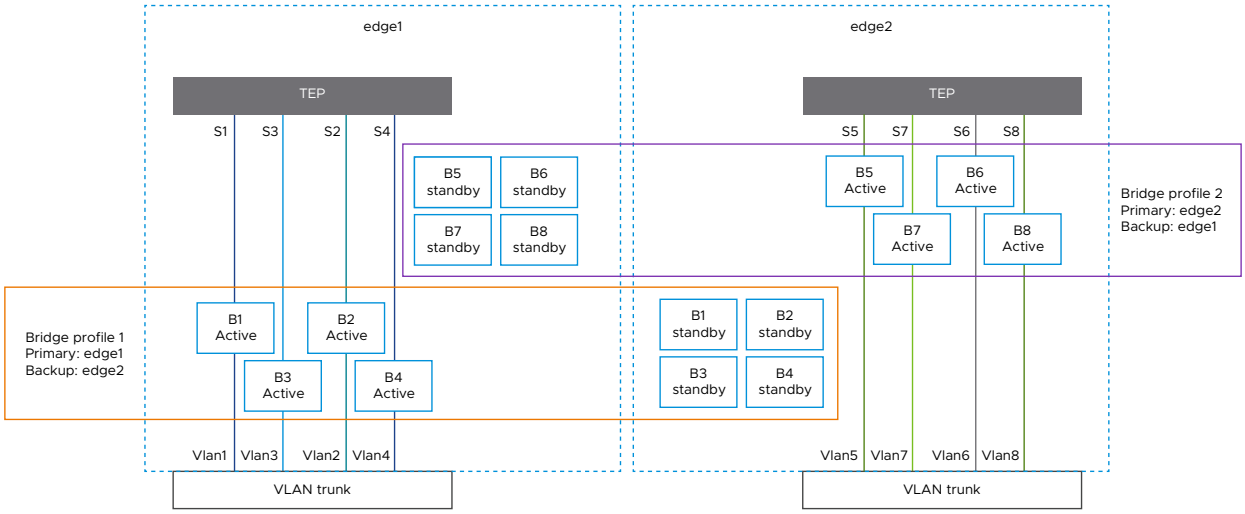
- Profile1 with edge1 as primary and edge2 as backup
- Profile2 with edge2 as primary and edge1 as backup

You can then create an arbitrary number of bridges using edge1 as primary (respectively backup), by associating them to the profile1 (respectively profile2). Those two profiles are enough to load share the bridged traffic between the two edges, on a per segment basis. The Few Bridge Profiles for Many Bridges diagram represents an example of bridging eight segments across two edges, using two edge bridge profiles.

This diagram shows bridge overlay segments S1 to VLAN 1, segment S2 to VLAN 2, and so on. Segments S1 to S4 are using bridge profile1, resulting in active bridges on edge1, standby on edge2. Segment S5 to S8 are using bridge profile2, leading to active bridges on edge2, standby on edge1. This diagram shows load sharing of the bridging functionality, on a per segment basis.

Select the Failover Mode

Figure 25-5. Few Bridge Profiles for Many Bridges



The benefit of the preemptive mode is that the system is attempting to forward the bridged traffic along a deterministic path. If you take the example of the Few Bridge Profiles for Many Bridges figure, with a preemptive mode, you are sure that the bridge traffic gets distributed on a per segment basis as soon as both edges are available, thus providing more bandwidth.

Depending on the availability of the edges and the failover mode selected for the bridge profiles, the active bridges might be running on the backup edges.

Select the Failover Mode

When both edges in the bridge profile are available, the active bridge is typically running on the primary edge. If the active bridge or the primary edge fails, the standby bridge on the backup edge takes over the active role and starts forwarding traffic between overlay segment and VLAN.

A bridge switchover, moving the active bridge to a different edge, is an operation that results in traffic loss. The bridge that is becoming active synchronizes the mac addresses that were learned on the previously active bridge and starts flooding RARP packets, using those mac addresses as source mac addresses. This mechanism is necessary to update the mac address tables of the physical infrastructure.

The Preemptive Mode

For example, what if a failure occurs on the primary edge and the bridge running on the backup edge is already active? In preemptive mode, when the failure is recovered on the primary edge, a bridge switchover is triggered and the bridge on the primary edge becomes active again.

The drawback of the preemptive mode is that there is a disruptive bridge convergence when the bridge on the primary edge recovers and becomes active again.

The Non-Preemptive Mode

In non-preemptive mode, the bridge on the primary edge recovers from failure as a standby bridge. The benefit of this mode is that there is no additional traffic disruption when the primary recovers. The preemptive mode is the best option in terms of bandwidth, thanks to its load sharing. The drawback of the non-preemptive mode is that bridge traffic flow is non-deterministic and can be sub-optimal. In the example shown in the Few Bridge Profiles for Many Bridges figure, after a failed edge recovers, the bridge traffic still flows through a unique edge, with no load sharing.

You can manually trigger a bridge switchover. To manually trigger a bridge switchover from the CLI of the edge currently hosting the standby bridge, enter: `set bridge <uuid> state active`.

Use this command only in non-preemptive mode. If you use it in preemptive mode, it returns an error.

Use this command only in non-preemptive mode. If you use it in preemptive mode, it returns an error.

For more information on set or get bridge commands, see the NSX-T Data Center Command-Line Interface Reference.

Prerequisites

- Verify that you have an NSX Edge cluster with two NSX Edge transport nodes.
- Verify that you are in Manager mode.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Go to Manager mode.
3. Select **Networking > Logical Switches > Edge Bridge Profiles > Add**.
4. Enter a name for the Edge bridge profile and optionally a description.
5. Select an NSX Edge cluster.
6. Select a primary node.
7. Select a backup node.
8. Select a failover node.

   The options are **Preemptive** and **Non-Preemptive**.
9. Click **Add**.

What to do next

Create a bridge-backed segment by extending an overlay segment to a VLAN or a range of VLANs.

Extend an Overlay Segment to a VLAN or a Range of VLANs in Manager Mode

After you have identified the edges on which you want the bridging functionality to be performed and created the appropriate edge bridge profile, the final step is to edit the segment configuration and specify the edge bridge profile to which you want to associate with the segment and the VLAN ID or range of VLAN IDs to which to bridge your segment. This will instantiate one or two bridges on the edges identified in the edge bridge profile.

When you configure a bridge with a single VLAN ID, a frame received on the overlay segment by the bridge gets decapsulated and forwarded on the VLAN uplink of the bridge with an added 802.1Q tag corresponding to this VLAN ID.

When you create the bridge specifying a VLAN ID range, you must configure the overlay segment being bridged for Guest VLAN Tagging (GVT). This means that the encapsulated frames already carry an 802.1Q tag. When the bridge receives an encapsulated frame carrying a VLAN tag on its overlay interface, it first checks that VLAN ID in the tag belongs to the VLAN range configured for the bridge. If this is the case, it forwards the frame on the VLAN uplink of the bridge carrying the original 802.1Q tag that was received on the overlay. Otherwise, it drops the frame.

Note If needed, you can configure multiple bridges on the same segment but:

- The same segment cannot be bridged twice on the same edge.
- The bridge does not have any loop detection or prevention. If you configure multiple bridges to the same bridging domain on the VLAN side it results in a permanent bridging loop.

Configuring a Bridge-Backed Segment

Prerequisites

- You have identified an overlay segment you want to bridge.
- You have an edge bridge profile specifying one or two edges attached to the overlay transport zone of your segment.
- If you are using edge VMs, you have checked the configuration requirements in Configure an Edge VM for Bridging.

Procedure

1. From a browser, log in with admin privileges to an NSX Manager.
2. Select **Networking > Logical Switches**.
3. Click the name of an overlay switch (traffic type: overlay).
4. Click **Attach**.
5. To attach to an Edge bridge profile,
   a. Select an Edge bridge profile.
   b. Select a transport zone.

c Enter a VLAN ID.

d Click **Save**.

6 Connect VMs to the logical switch if they are not already connected.

The VMs must be on transport nodes in the same transport zone as the Edge bridge profile.

Logical Routers in Manager Mode

NSX supports a 2-tier routing model.

In the top tier is the tier-0 logical router. Northbound, the tier-0 logical router connects to one or more physical routers or layer 3 switches and serves as a gateway to the physical infrastructure. Southbound, the tier-0 logical router connects to one or more tier-1 logical routers or directly to one or more logical switches.

In the bottom tier is the tier-1 logical router. Northbound, the tier-1 logical router connects to a tier-0 logical router. Southbound, it connects to one or more logical switches.

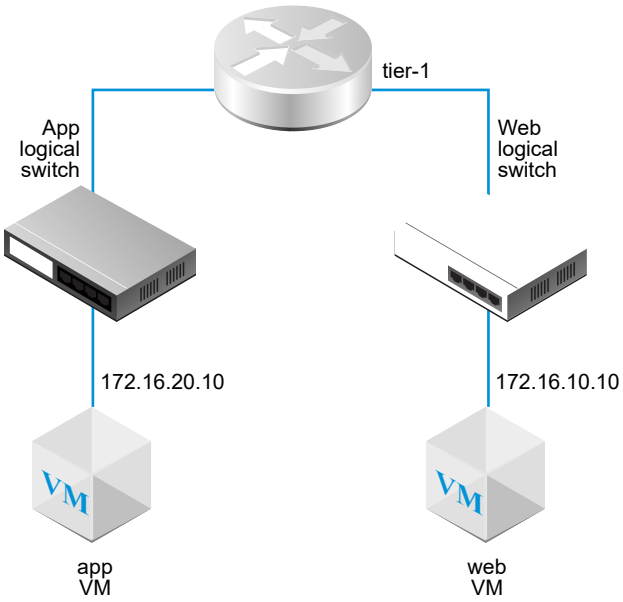
Note If you use **Manager** mode to modify objects created in the **Policy** mode, some settings might not be configurable. These read-only settings have this icon next to them: ⊖. See Chapter 1 NSX Manager for more information.

Tier-1 Logical Router

Tier-1 logical routers have downlink ports to connect to logical switches and uplink ports to connect to tier-0 logical routers.

When you add a logical router, it is important that you plan the networking topology you are building.

Figure 25-6. Tier-1 Logical Router Topology



For example, this simple topology shows two logical switches connected to a tier-1 logical router. Each logical switch has a single VM connected. The two VMs can be on different hosts or the same host, in different host clusters or in the same host cluster. If a logical router does not separate the VMs, the underlying IP addresses configured on the VMs must be in the same subnet. If a logical router does separate them, the IP addresses on the VMs must be in different subnets.

In some scenarios, external clients send ARP queries for MAC addresses bound to LB VIP ports. However, LB VIP ports do not have MAC addresses and cannot handle such queries. Proxy ARP is implemented on the centralized service ports of a tier-1 logical router to handle ARP queries on behalf of the LB VIP ports.

When a tier-1 logical router is configured with DNAT, Edge firewall, and load balancer, traffic to and from another tier-1 logical router is processed in this order: DNAT first, then Edge firewall, and then load balancer. Traffic within the tier-1 logical router is processed through DNAT first and then load balancer. Edge firewall processing is skipped.

On a tier-0 or tier-1 logical router, you can configure different types of ports. One type is called centralized service port (CSP). You must configure a CSP on a tier-0 logical router in active-standby mode or a tier-1 logical router to connect to a VLAN-backed logical switch, or to create a standalone tier-1 logical router. A CSP supports the following services on a tier-0 logical router in active-standby mode or a tier-1 logical router:

- NAT
- Load balancing
- Stateful firewall
- VPN (IPsec and L2VPN)

Create a Tier-1 Logical Router in Manager Mode

The tier-1 logical router must be connected to the tier-0 logical router to get the northbound physical router access.

Prerequisites

- Verify that the logical switches are configured. See Create a Logical Switch in Manager Mode.
- Verify that an NSX Edge cluster is deployed to perform network address translation (NAT) configuration. See the *NSX Installation Guide*.
- Familiarize yourself with the tier-1 logical router topology. See Tier-1 Logical Router.
- Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Select **Networking > Tier-1 Logical Routers > Add**.
3. Enter a name for the logical router and optionally a description.
4. (Optional) Select a tier-0 logical router to connect to this tier-1 logical router.

   If you do not yet have any tier-0 logical routers configured, you can leave this field blank for now and edit the router configuration later.
5. (Optional) Select an NSX Edge cluster.

   To deselect a cluster that you selected, click the **x** icon. If the tier-1 logical router is going to be used for NAT configuration, it must be connected to an NSX Edge cluster. If you do not yet have any NSX Edge clusters configured, you can leave this field blank for now and edit the router configuration later.
6. (Optional) Click the **StandBy Relocation** toggle to enable or disable standby relocation.

   Standby relocation means that if the Edge node where the active or standby logical router is running fails, a new standby logical router is created on another Edge node to maintain high availability. If the Edge node that fails is running the active logical router, the original standby logical router becomes the active logical router and a new standby logical router is created. If the Edge node that fails is running the standby logical router, the new standby logical router replaces it.

7 (Optional) If you selected an NSX Edge cluster, select a failover mode.

Option	Description
Preemptive	If the preferred node fails and recovers, it will preempt its peer and become the active node. The peer will change its state to standby. This is the default option.
Non-preemptive	If the preferred node fails and recovers, it will check if its peer is the active node. If so, the preferred node will not preempt its peer and will be the standby node.

8 (Optional) Click the **Advanced** tab and enter a value for **Intra Tier-1 Transit Subnet**.

9 Click **Add**.

Results

After the logical router is created, if you want to remove the Edge cluster from the router's configuration, perform the following steps:

- Click the name of the router to see the configuration details.
- Select **Services** > **Edge Firewall**.
- Click **Disable Firewall**.
- Click the **Overview** tab and click **Edit**.
- In the **Edge Cluster** field, click the **x** icon.
- Click **Save**.

If this logical router supports more than 5000 VMs, you must run the following commands on each node of the NSX Edge cluster to increase the size of the ARP table.

```
set debug-mode
set dataplane neighbor max-arp-logical-router 10000
```

You must re-run the commands after a dataplane restart or a node reboot because the change is not persistent.

What to do next

Create downlink ports for your tier-1 logical router. See Add a Downlink Port on a Tier-1 Logical Router in Manager Mode.

Add a Downlink Port on a Tier-1 Logical Router in Manager Mode

When you create a downlink port on a tier-1 logical router, the port serves as a default gateway for the VMs that are in the same subnet.

Prerequisites

- Verify that a tier-1 logical router is configured. See Create a Tier-1 Logical Router in Manager Mode .

- Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Networking** > **Tier-1 Logical Routers**.

3 Click the name of a tier-1 router.

4 Click the **Configuration** tab and select **Router Ports**.

5 Click **Add**.

6 Enter a name for the router port and optionally a description.

7 In the **Type** field, select **Downlink**.

8 For **URPF Mode**, select **Strict** or **None**.

URPF (unicast Reverse Path Forwarding) is a security feature.

9 (Optional) Select a logical switch.

10 Select whether this attachment creates a switch port or updates an existing switch port.

If the attachment is for an existing switch port, select the port from the drop-down menu.

11 Enter an IP address and a prefix length for the router port.

12 (Optional) Select a DHCP relay service.

13 Click **Add**.

What to do next

Enable route advertisement to provide North-South connectivity between VMs and external physical networks or between different tier-1 logical routers that are connected to the same tier-0 logical router. See Configure Route Advertisement on a Tier-1 Logical Router in Manager Mode.

Add a VLAN Port on a Tier-0 or Tier-1 Logical Router in Manager Mode

If you have only VLAN-backed logical switches, you can connect the switches to VLAN ports on a tier-0 or tier-1 router so that NSX can provide layer-3 services.

Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Locate the router in **Networking > Tier-0 Logical Routers** or **Networking > Tier-1 Logical Routers** and select it.

3 Click the **Configuration** tab and select **Router Ports**.

4 Click **Add**.

5 Enter a name for the router port and optionally a description.

6 In the **Type** field, select **Centralized**.

7 For **URPF Mode**, select **Strict** or **None**.

URPF (unicast Reverse Path Forwarding) is a security feature.

8 (Required) Select a logical switch.

9 Select whether this attachment creates a switch port or updates an existing switch port.

If the attachment is for an existing switch port, select the port from the drop-down menu.

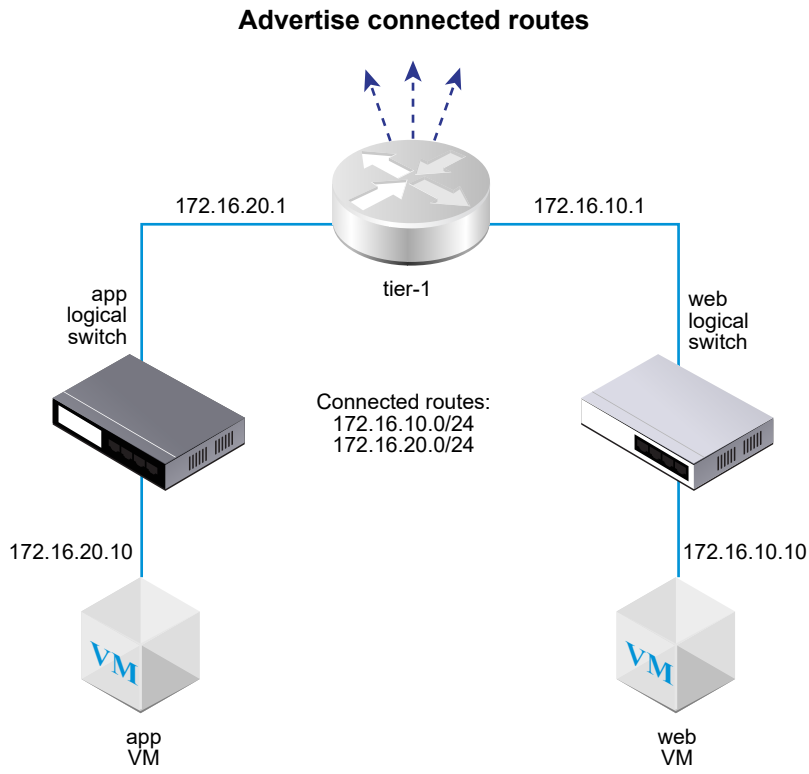
10 Enter the router port IP address in CIDR notation.

11 Click **Add**.

Configure Route Advertisement on a Tier-1 Logical Router in Manager Mode

To provide Layer 3 connectivity between VMs connected to logical switches that are attached to different tier-1 logical routers, it is necessary to enable tier-1 route advertisement towards tier-0. You do not need to configure a routing protocol or static routes between tier-1 and tier-0 logical routers. NSX creates NSX static routes automatically when you enable route advertisement.

For example, to provide connectivity to and from the VMs through other peer routers, the tier-1 logical router must have route advertisement configured for connected routes. If you don't want to advertise all connected routes, you can specify which routes to advertise.



Prerequisites

- Verify that VMs are attached to logical switches. See Logical Switches in Manager Mode.
- Verify that downlink ports for the tier-1 logical router are configured. See Add a Downlink Port on a Tier-1 Logical Router in Manager Mode.
- Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Networking > Tier-1 Logical Routers**.

3 Click the name of a tier-1 router.

4 Select **Route Advertisement** from the **Routing** drop-down menu.

5 Click **Edit** to edit the route advertisement configuration.

You can toggle the following switches:

- **Status**
- **Advertise All NSX Connected Routes**
- **Advertise All NAT Routes**

- **Advertise All Static Routes**
- **Advertise All LB VIP Routes**
- **Advertise All LB SNAT IP Routes**
- **Advertise All DNS Forwarder Routes**

a Click **Save**.

6 Click **Add** to advertise routes.

a Enter a name and optionally a description.

b Enter a route prefix in CIDR format.

c Click **Apply Filter** to set the following options:

Action	Specify **Allow** or **Deny**.
Match route types	Select one or more of the following: ■ **Any** ■ **NSX Connected** ■ **Tier-1 LB VIP** ■ **Static** ■ **Tier-1 NAT** ■ **Tier-1 LB SNAT**
Prefix operator	Select **GE** (greater than or equal) or **EQ** (equal).

d Click **Add**.

What to do next

Familiarize yourself with the tier-0 logical router topology and create the tier-0 logical router. See Tier-0 Logical Router.

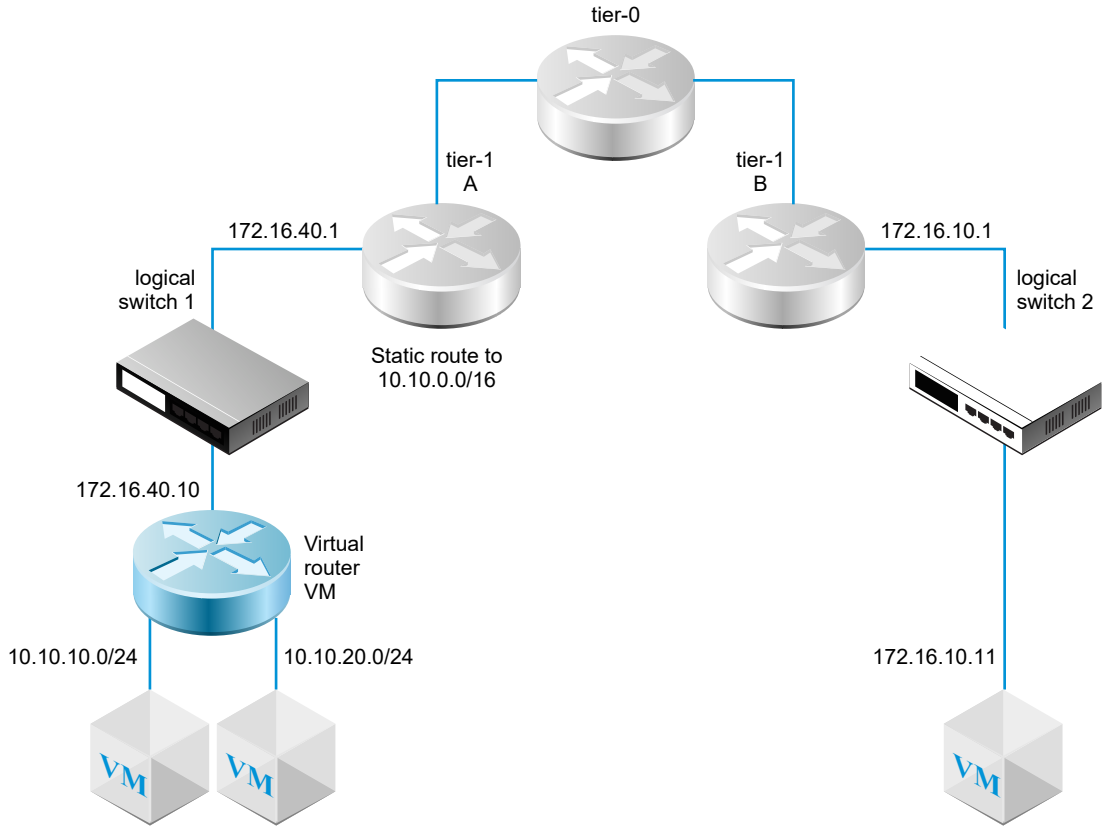
If you already have a tier-0 logical router connected to the tier-1 logical router, you can verify that the tier-0 router is learning the tier-1 router connected routes. See Verify that a Tier-0 Router Has Learned Routes from a Tier-1 Router.

Configure a Tier-1 Logical Router Static Route in Manager Mode

You can configure a static route on a tier-1 logical router to provide connectivity from NSX to a set of networks that are accessible through a virtual router.

For example, in the following diagram, the tier-1 A logical router has a downlink port to an NSX logical switch. This downlink port (172.16.40.1) serves the default gateway for the virtual router VM. The virtual router VM and tier-1 A are connected through the same NSX logical switch. The tier-1 logical router has a static route 10.10.0.0/16 that summarizes the networks available through the virtual router. Tier-1 A then has route advertisement configured to advertise the static route to tier-1 B.

Figure 25-7. Tier-1 Logical Router Static Route Topology



Recursive static routes are supported.

Prerequisites

- Verify that a downlink port is configured. See Add a Downlink Port on a Tier-1 Logical Router in Manager Mode.
- Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Verify that a downlink port is configured. See Add a Downlink Port on a Tier-1 Logical Router in Manager Mode.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Select **Networking > Tier-1 Logical Routers**.
3. Click the name of a tier-1 router.
4. Click the **Routing** tab and select **Static Routes** from the drop-down menu.
5. Click **Add**.

6 Enter a network address in the CIDR format.

   Static route based on IPv6 is supported. IPv6 prefixes can only have an IPv6 next hop.

   For example, 10.10.10.0/16 or an IPv6 address.

7 Click **Add** to add a next-hop IP address.

   For example, 172.16.40.10. You can also specify a null route by clicking the pencil icon and selecting **NULL** from the drop-down. To add another next hop addresses, click **Add** again.

8 Click **Add** at the bottom of the dialog box.

   The newly created static route network address appears in the row.

9 From the tier-1 logical router, select **Routing > Route Advertisement**.

10 Click **Edit** and select **Advertise All Static Routes**.

11 Click **Save**.

   The static route is propagated across the NSX overlay.

Create a Standalone Tier-1 Logical Router in Manager Mode

A standalone tier-1 logical router has no downlink and no connection to a tier-0 router. It has a service router but no distributed router. The service router can be deployed on one NSX Edge node or two NSX Edge nodes in active-standby mode.

A standalone tier-1 logical router:

- Must not have a connection to a tier-0 logical router.
- Can have only one centralized service port (CSP) if it is used to attach a load balancer (LB) service.
- Can connect to an overlay logical switch or a VLAN logical switch.
- Supports any combination of the services IPSec, NAT, firewall, load balancer, and service insertion. For ingress, the order of processing is: IPSec – DNAT – firewall – load balancer - service insertion. For egress, the order of processing is: service insertion - load balancer - firewall - SNAT - IPSec.

Typically, a standalone tier-1 logical router is connected to a logical switch that a regular tier-1 logical router is also connected to. The standalone tier-1 logical router can communicate with other devices through the regular tier-1 logical router after static routes and route advertisements are configured.

Before using the standalone tier-1 logical router, note the following:

- To specify the default gateway for the standalone tier-1 logical router, you must add a static route. The subnet should be 0.0.0.0/0 and the next hop is the IP address of a regular tier-1 router connected to the same switch.

- ARP proxy on the standalone router is supported. You can configure an LB virtual server IP or LB SNAT IP in the CSP's subnet. For example, if the CSP IP is 1.1.1.1/24, the virtual IP can be 1.1.1.2. It can also be an IP in another subnet such as 2.2.2.2 if routing is properly configured so that traffic for 2.2.2.2 can reach the standalone router.
- For an NSX Edge VM, you cannot have more than one CSPs which are connected to the same VLAN-backed logical switch or different VLAN-backed logical switches that have the same VLAN ID.

Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Networking > Tier-1 Logical Routers > Add**.

3 Enter a name for the logical router, and optionally a description.

4 (Required) Select an NSX Edge cluster to connect to this tier-1 logical router.

5 (Required) Select a failover mode and cluster members.

Option	Description
Preemptive	If the preferred node fails and recovers, it will preempt its peer and become the active node. The peer will change its state to standby. This is the default option.
Non-preemptive	If the preferred node fails and recovers, it will check if its peer is the active node. If so, the preferred node will not preempt its peer and will be the standby node.

6 Click **Add**.

7 Click the name of the router that you just created.

8 Click the **Configuration** tab and select **Router Ports**.

9 Click **Add**.

10 Enter a name for the router port and optionally a description.

11 In the **Type** field, select **Centralized**.

12 For **URPF Mode**, select **Strict** or **None**.

URPF (Unicast Reverse Path Forwarding) is a security feature.

13 (Required) Select a logical switch.

14 Select whether this attachment creates a switch port or updates an existing switch port.

15 Enter the router port IP address in CIDR notation.

16 Click **Add**.

Tier-0 Logical Router

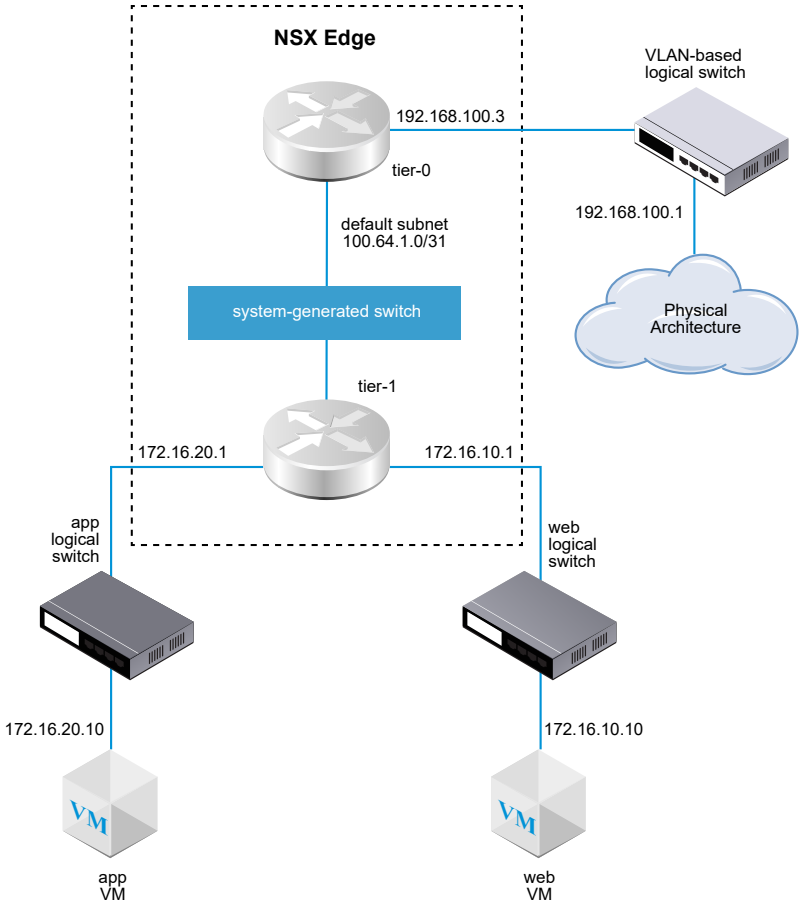
A tier-0 logical router provides a gateway service between the logical and physical network.

NSX Cloud Note If using NSX Cloud, see NSX Features Supported with NSX Cloud for a list of auto-generated logical entities, supported features, and configurations required for NSX Cloud.

An Edge node can support only one tier-0 gateway or logical router. When you create a tier-0 gateway or logical router, make sure you do not create more tier-0 gateways or logical routers than the number of Edge nodes in the NSX Edge cluster.

When you add a tier-0 logical router, it is important that you map out the networking topology you are building.

Figure 25-8. Tier-0 Logical Router Topology



For simplicity, the sample topology shows a single tier-1 logical router connected to a single tier-0 logical router hosted on a single NSX Edge node. Keep in mind that this is not a recommended topology. Ideally, you should have a minimum of two NSX Edge nodes to take full advantage of the logical router design.

The tier-1 logical router has a web logical switch and an app logical switch with respective VMs attached. The router-link switch between the tier-1 router and the tier-0 router is created automatically when you attach the tier-1 router to the tier-0 router. Thus, this switch is labeled as system generated.

In some scenarios, external clients send ARP queries for MAC addresses bound to loopback or IKE IP ports. However, loopback and IKE IP ports do not have MAC addresses and cannot handle such queries. Proxy ARP is implemented on the uplink and centralized service ports of a tier-0 logical router to handle ARP queries on behalf of the loopback and IKE IP ports.

When a tier-0 logical router is configured with DNAT, IPsec, and Edge firewall, traffic is processed in this order: IPsec first, then DNAT, and then Edge firewall.

On a tier-0 or tier-1 logical router, you can configure different types of ports. One type is called centralized service port (CSP). You must configure a CSP on a tier-0 logical router in active-standby mode or a tier-1 logical router to connect to a VLAN-backed logical switch, or to create a standalone tier-1 logical router. A CSP supports the following services on a tier-0 logical router in active-standby mode or a tier-1 logical router:

- NAT
- Load balancing
- Stateful firewall
- VPN (IPsec and L2VPN)

Create a Tier-0 Logical Router in Manager Mode

Tier-0 logical routers have downlink ports to connect to NSX tier-1 logical routers and uplink ports to connect to external networks.

Prerequisites

- Verify that at least one NSX Edge is installed. See the *NSX Installation Guide*
- Verify that an NSX Edge cluster is configured. See the *NSX Installation Guide*.
- Familiarize yourself with the networking topology of the tier-0 logical router. See Tier-0 Logical Router.
- Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Select **Networking > Tier-0 Logical Routers > Add**.
3. Enter a name for the tier-0 logical router.

4 Select an existing NSX Edge cluster from the drop-down menu to back this tier-0 logical router.

5 (Optional) Select a high-availability mode.

By default, the active-active mode is used. In the active-active mode, traffic is load balanced across all members. In active-standby mode, all traffic is processed by an elected active member. If the active member fails, a new member is elected to be active.

6

7 (Optional) Click the **Advanced** tab to enter a subnet for the intra-tier 0 transit subnet.

This is the subnet that connects to the tier-0 services router to its distributed router. If you leave this blank, the default 169.0.0.0/28 subnet is used.

8 (Optional) Click the **Advanced** tab to enter a subnet for the tier-0-tier-1 transit subnet.

This is the subnet that connects the tier-0 router to any tier-1 routers that connect to this tier-0 router. If you leave this blank, the default address space assigned for these tier-0-to-tier-1 connections is 100.64.0.0/16. Each tier-0-to-tier-1 peer connection is provided a /31 subnet within the 100.64.0.0/16 address space.

9 Click **Save**.

The new tier-0 logical router appears as a link.

10 (Optional) Click the tier-0 logical router link to review the summary.

What to do next

Attach tier-1 logical routers to this tier-0 logical router.

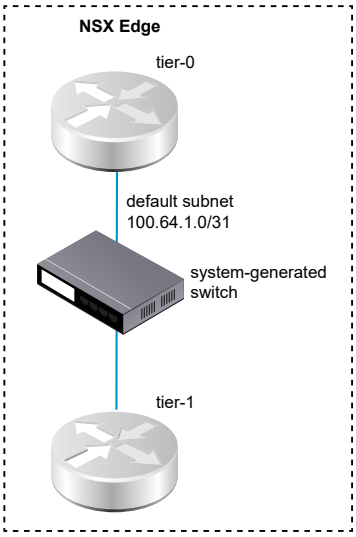
Configure the tier-0 logical router to connect it to a VLAN logical switch to create an uplink to an external network. See Connect a Tier-0 Logical Router to a VLAN Logical Switch for the NSX Edge Uplink in Manager Mode.

Attach Tier-1 Router to a Tier-0 Router in Manager Mode

You can attach the tier-0 logical router to the tier-1 logical router so that the tier-1 logical router gets northbound and east-west network connectivity.

When you attach a tier-1 logical router to a tier-0 logical router, a router-link switch between the two routers is created. This switch is labeled as system-generated in the topology. The default address space assigned for these tier-0-to-tier-1 connections is 100.64.0.0/16. Each tier-0-to-tier-1 peer connection is provided a /31 subnet within the 100.64.0.0/16 address space. Optionally, you can configure the address space in the tier-0 **Summary > Advanced** configuration.

The following figure shows a sample topology.



Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Select **Networking > Tier-1 Logical Routers**.
3. Select the tier-1 logical router.
4. In the Tier-0 Connection section, click **Connect**.
5. Select a tier-0 logical router from the drop-down menu.
6. (Optional) Select an NSX Edge cluster from the drop-down menu.

   The tier-1 router needs to be backed by an edge device if the router is going to be used for services, such as NAT. If you do not select an NSX Edge cluster, the tier-1 router cannot perform NAT.
7. Specify members and a preferred member.

   If you select an NSX Edge cluster and leave the members and preferred member fields blank, NSX sets the backing edge device from the specified cluster for you.
8. Click **Save**.
9. Click the **Configuration** tab of the tier-1 router to verify that a new point-to-point linked port IP address is created.

   For example, the IP address of the linked port can be 100.64.1.1/31.
10. Select the tier-0 logical router from the navigation panel.

11 Click the **Configuration** tab of the tier-0 router to verify that a new point-to-point linked port IP address is created.

   For example, the IP address of the linked port can be 100.64.1.1/31.

What to do next

Verify that the tier-0 router is learning routes that are advertised by the tier-1 routers.

Verify that a Tier-0 Router Has Learned Routes from a Tier-1 Router

When a tier-1 logical router advertises routes to a tier-0 logical router, the routes are listed in the tier-0 router's routing table as NSX static routes.

Procedure

1 On the NSX Edge, run the `get logical-routers` command to find the VRF number of the tier-0 service router.

```
nsx-edge-1> get logical-routers
Logical Router
UUID        : 736a80e3-23f6-5a2d-81d6-bbefb2786666
vrf         : 0
type        : TUNNEL

Logical Router
UUID        : 421a2d0d-f423-46f1-93a1-2f9e366176c8
vrf         : 5
type        : SERVICE_ROUTER_TIER0

Logical Router
UUID        : f3ce9d7d-7123-47d6-aba6-45cf1388ca7b
vrf         : 6
type        : DISTRIBUTED_ROUTER

Logical Router
UUID        : c8e64eff-02b2-4462-94ff-89f3788f1a61
vrf         : 7
type        : SERVICE_ROUTER_TIER1

Logical Router
UUID        : fb6c3f1f-599f-4421-af8a-99692dff3dd4
vrf         : 8
type        : DISTRIBUTED_ROUTER
```

2 Run the `vrf <number>` command to enter the tier-0 service router context.

```
nsx-edge-1> vrf 5
nsx-edge1(tier0_sr)>
```

3 On the tier-0 service router, run the `get route` command and make sure the expected routes appear in the routing table.

Notice that the NSX static routes (ns) are learned by the tier-0 router because the tier-1 router is advertising routes.

```
nsx-edge1(tier0_sr)> get route

Flags: c - connected, s - static, b - BGP, ns - nsx_static
nc - nsx_connected, rl - router_link, t0n: Tier0-NAT, t1n: Tier1-NAT

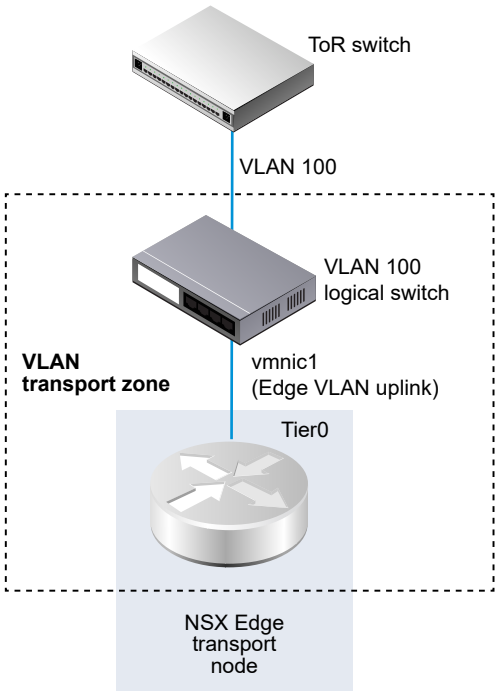
Total number of routes: 7

b    10.10.10.0/24        [20/0]        via 192.168.100.254
rl   100.91.176.0/31      [0/0]         via 169.254.0.1
c    169.254.0.0/28       [0/0]         via 169.254.0.2
ns   172.16.10.0/24       [3/3]         via 169.254.0.1
ns   172.16.20.0/24       [3/3]         via 169.254.0.1
c    192.168.100.0/24     [0/0]         via 192.168.100.2
```

Connect a Tier-0 Logical Router to a VLAN Logical Switch for the NSX Edge Uplink in Manager Mode

To create an NSX Edge uplink, you must connect a tier-0 router to a VLAN switch.

The following simple topology shows a VLAN logical switch inside of a VLAN transport zone. The VLAN logical switch has a VLAN ID that matches the VLAN ID on the TOR port for the Edge's VLAN uplink.



Prerequisites

- Create a VLAN logical switch. See Create a VLAN Logical Switch for the NSX Edge Uplink in Manager Mode.
- Create a tier-0 router.
- Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Select **Networking > Tier-0 Logical Routers**.
3. Select the tier-0 logical router.
4. From the **Configuration** tab, add a new logical router port.
5. Type a name for the port, such as uplink.
6. Select the **Uplink** type.
7. Select an edge transport node.
8. Select a VLAN logical switch.
9. Select either **Attach to new switch port** or **Attach to existing switch port**.

   If you select **Attach to new switch port,** the port will be automatically created.
10. Specify an IP address that is in the same subnet as the connected port on the TOR switch.

Results

A new uplink port is added for the tier-0 router.

What to do next

Configure BGP or a static route.

Verify the Tier-0 Logical Router and TOR Connection

For routing to work on the uplink from the tier-0 router, connectivity with the top-of-rack device must be in place.

Prerequisites

- Verify that the tier-0 logical router is connected to a VLAN logical switch. See Connect a Tier-0 Logical Router to a VLAN Logical Switch for the NSX Edge Uplink in Manager Mode.

Procedure

1. Log in to the NSX Edge CLI.

2 On the NSX Edge, run the `get logical-routers` command to find the VRF number of the tier-0 service router.

```
nsx-edge-1> get logical-routers
Logical Router
UUID        : 736a80e3-23f6-5a2d-81d6-bbefb2786666
vrf         : 0
type        : TUNNEL

Logical Router
UUID        : 421a2d0d-f423-46f1-93a1-2f9e366176c8
vrf         : 5
type        : SERVICE_ROUTER_TIER0

Logical Router
UUID        : f3ce9d7d-7123-47d6-aba6-45cf1388ca7b
vrf         : 6
type        : DISTRIBUTED_ROUTER

Logical Router
UUID        : c8e64eff-02b2-4462-94ff-89f3788f1a61
vrf         : 7
type        : SERVICE_ROUTER_TIER1

Logical Router
UUID        : fb6c3f1f-599f-4421-af8a-99692dff3dd4
vrf         : 8
type        : DISTRIBUTED_ROUTER
```

3 Run the `vrf <number>` command to enter the tier-0 service router context.

```
nsx-edge-1> vrf 5
nsx-edge1(tier0_sr)>
```

4 On the tier-0 service router, run the `get route` command and make sure the expected route appears in the routing table.

Notice that the route to the TOR appears as connected (c).

```
nsx-edge1(tier0_sr)> get route
Flags: t0c - Tier0-Connected, t0s - Tier0-Static, b - BGP,
t0n - Tier0-NAT, t1s - Tier1-Static, t1c - Tier1-Connected,
t1n: Tier1-NAT, t1l: Tier1-LB VIP, t1ls: Tier1-LB SNAT,
t1d: Tier1-DNS FORWARDER, t1ipsec: Tier1-IPSec, isr: Inter-SR,
> - selected route, * - FIB route

Total number of routes: 11

t1c> * 1.1.1.0/25 [3/0] via 100.64.1.1, downlink-282, 08w4d03h
t1c> * 1.1.2.0/24 [3/0] via 100.64.1.1, downlink-282, 08w4d03h
```

```
t0c> * 1.1.3.0/24 is directly connected, downlink-275, 08w4d03h
b  > * 2.1.4.0/24 [20/0] via 40.40.40.10, uplink-273, 01w0d02h
b  > * 10.182.48.0/20 [20/0] via 40.40.40.10, uplink-273, 01w0d02h
t0c> * 40.40.40.0/24 is directly connected, uplink-273, 08w4d03h
t0c> * 100.64.1.0/31 is directly connected, downlink-282, 08w4d03h
t0c> * 169.254.0.0/24 is directly connected, downlink-277, 01w0d02h
b  > * 172.17.0.0/16 [20/0] via 40.40.40.10, uplink-273, 01w0d02h
t0c> * fc36:a750:db0d:7800::/64 is directly connected, downlink-282, 08w4d03h
t0c> * fe80::/64 is directly connected, downlink-282, 08w4d03h
```

5 Ping the TOR.

```
nsx-edge1(tier0_sr)> ping    192.168.100.254
PING 192.168.100.254 (192.168.100.254): 56 data bytes
64 bytes from 192.168.100.254: icmp_seq=0 ttl=64 time=2.822 ms
64 bytes from 192.168.100.254: icmp_seq=1 ttl=64 time=1.393 ms
^C
nsx-edge1>
--- 192.168.100.254 ping statistics ---
3 packets transmitted, 2 packets received, 33.3% packet loss
round-trip min/avg/max/stddev = 1.393/2.107/2.822/0.715 ms
```

Results

Packets are sent between the tier-0 logical router and physical router to verify a connection.

What to do next

Depending on your networking requirements, you can configure a static route or BGP. See Configure a Static Route in Manager Mode or Configure BGP on a Tier-0 Logical Router in Manager Mode.

Add a Loopback Router Port in Manager Mode

You can add a loopback port to a tier-0 logical router.

The loopback port can be used for the following purposes:

- Router ID for routing protocols
- NAT
- BFD
- Source address for routing protocols

Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Networking > Tier-0 Logical Routers**.

3 Select the tier-0 logical router.

4 Select **Configuration > Router Ports**

5 Click **Add**.

6 Enter a name and optionally a description.

7 Select the **Loopback** type.

8 Select an edge transport node.

9 Enter an IP address in CIDR format.

Results

A new port is added for the tier-0 router.

Add a VLAN Port on a Tier-0 or Tier-1 Logical Router in Manager Mode

If you have only VLAN-backed logical switches, you can connect the switches to VLAN ports on a tier-0 or tier-1 router so that NSX can provide layer-3 services.

Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Locate the router in **Networking > Tier-0 Logical Routers** or **Networking > Tier-1 Logical Routers** and select it.

3 Click the **Configuration** tab and select **Router Ports**.

4 Click **Add**.

5 Enter a name for the router port and optionally a description.

6 In the **Type** field, select **Centralized**.

7 For **URPF Mode**, select **Strict** or **None**.

URPF (unicast Reverse Path Forwarding) is a security feature.

8 (Required) Select a logical switch.

9 Select whether this attachment creates a switch port or updates an existing switch port.

If the attachment is for an existing switch port, select the port from the drop-down menu.

10 Enter the router port IP address in CIDR notation.

11 Click **Add**.

Configure High Availability VIP in Manager Mode

With HA VIP (high availability virtual IP) configured, a tier-0 logical router is operational even if one uplink is down. The physical router interacts with the HA VIP only.

Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Networking > Tier-0 Logical Routers**.

3 Click the tier-0 logical router name.

4 Click **Configuration > HA VIP**.

5 Click **Add**.

6 Enter an IP address in CIDR format.

7 To enable HA VIP, set the status to **Enabled**.

8 Select exactly two uplink ports.

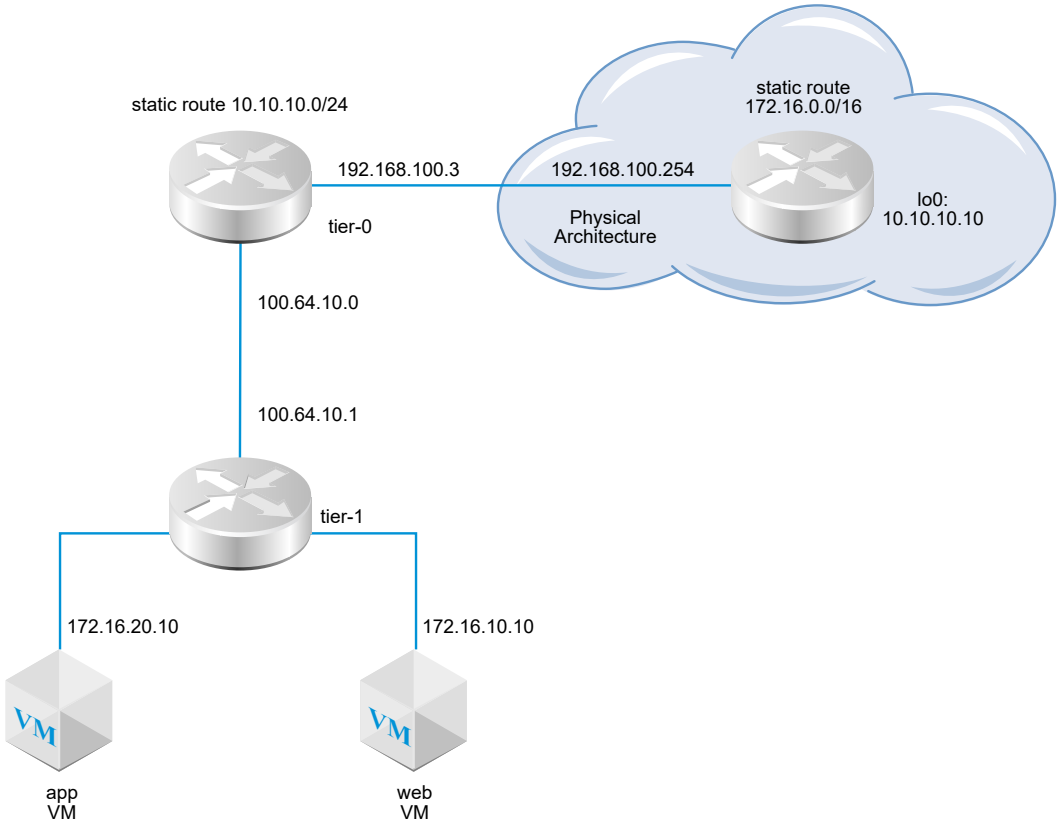
9 Click **Add**.

Configure a Static Route in Manager Mode

You can configure a static route on the tier-0 router to external networks. After you configure a static route, there is no need to advertise the route from tier-0 to tier-1, because tier-1 routers automatically have a static default route towards their connected tier-0 router.

The static route topology shows a tier-0 logical router with a static route to the 10.10.10.0/24 prefix in the physical architecture. For test purposes, the 10.10.10.10/32 address is configured on the external router loopback interface. The external router has a static route to the 172.16.0.0/16 prefix to reach the app and web VMs.

Figure 25-9. Static Route Topology



Recursive static routes are supported.

Prerequisites

- Verify that the physical router and tier-0 logical router are connected. See Verify the Tier-0 Logical Router and TOR Connection.
- Verify that the tier-1 router is configured to advertise connected routes. See Create a Tier-1 Logical Router in Manager Mode .
- Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Select **Networking > Tier-0 Logical Routers**.
3. Select the tier-0 logical router.
4. Click the **Routing** tab and select **Static Route** from the drop-down menu.
5. Select **Add**.

6 Enter a network address in the CIDR format.

   For example, 10.10.10.0/24.

7 Click **+ Add** to add a next-hop IP address.

   For example, 192.168.100.254. You can also specify a null route by clicking the pencil icon and selecting **NULL** from the drop-down.

8 Specify the administrative distance.

9 Select a logical router port from the dropdown list.

   The list includes IPSec Virtual Tunnel Interface (VTI) ports.

10 Click the **Add** button.

What to do next

Check that the static route is configured properly. See Verify the Static Route on a Tier-0 Router.

Verify the Static Route on a Tier-0 Router

Use the CLI to verify that the static route is connected. You must also verify the external router can ping the internal VMs and the internal VMs can ping the external router.

Prerequisites

Verify that a static route is configured. See Configure a Static Route in Manager Mode.

Procedure

1 Log in to the NSX Manager CLI.

2 Confirm the static route.

a Get the service router UUID information.

```
get logical-routers
```

```
nsx-edge1> get logical-routers
Logical Router
UUID        : 736a80e3-23f6-5a2d-81d6-bbefb2786666
vrf         : 2
type        : TUNNEL

Logical Router
UUID        : d40bbfa4-3e3d-4178-8615-6f42ea335037
vrf         : 4
type        : SERVICE_ROUTER_TIER0

Logical Router
UUID        : d0289ba4-250e-41b4-8ffc-7cab4a46c3e4
vrf         : 5
type        : DISTRIBUTED_ROUTER

Logical Router
UUID        : a6ee6316-2212-4171-99cc-930c98bcad7f
vrf         : 6
type        : DISTRIBUTED_ROUTER
```

b Locate the UUID information from the output.

```
Logical Router
UUID        : d40bbfa4-3e3d-4178-8615-6f42ea335037
vrf         : 4
type        : SERVICE_ROUTER_TIER0
```

c Verify that the static route works.

```
get logical-router d40bbfa4-3e3d-4178-8615-6f42ea335037 route static
```

```
Flags: c - connected, s - static, b - BGP, ns - nsx_static
nc - nsx_connected, rl - router_link, t0n: Tier0-NAT, t1n: Tier1-NAT

s    10.10.10.0/24        [1/1]         via 192.168.100.254
rl   100.64.1.0/31        [0/0]         via 169.0.0.1
ns   172.16.10.0/24       [3/3]         via 169.0.0.1
ns   172.16.20.0/24       [3/3]         via 169.0.0.1
```

3 From the external router, ping the internal VMs to confirm that they are reachable through the NSX overlay.

a Connect to the external router.

```
ping 172.16.10.10
```

```
PING 172.16.10.10 (172.16.10.10) 56(84) bytes of data.
64 bytes from 172.16.10.10: icmp_req=1 ttl=62 time=127 ms
64 bytes from 172.16.10.10: icmp_req=2 ttl=62 time=1.96 ms
^C
--- 172.16.10.10 ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 1001ms
rtt min/avg/max/mdev = 1.966/64.793/127.620/62.827 ms
```

b Test the network connectivity.

```
traceroute 172.16.10.10
```

```
traceroute to 172.16.10.10 (172.16.10.10), 30 hops max, 60 byte packets
 1  192.168.100.3 (192.168.100.3)  0.640 ms  0.575 ms  0.696 ms
 2  100.64.1.1 (100.64.1.1)  0.656 ms  0.604 ms  0.578 ms
 3  172.16.10.10 (172.16.10.10)  3.397 ms  3.703 ms  3.790 ms
```

4 From the VMs, ping the external IP address.

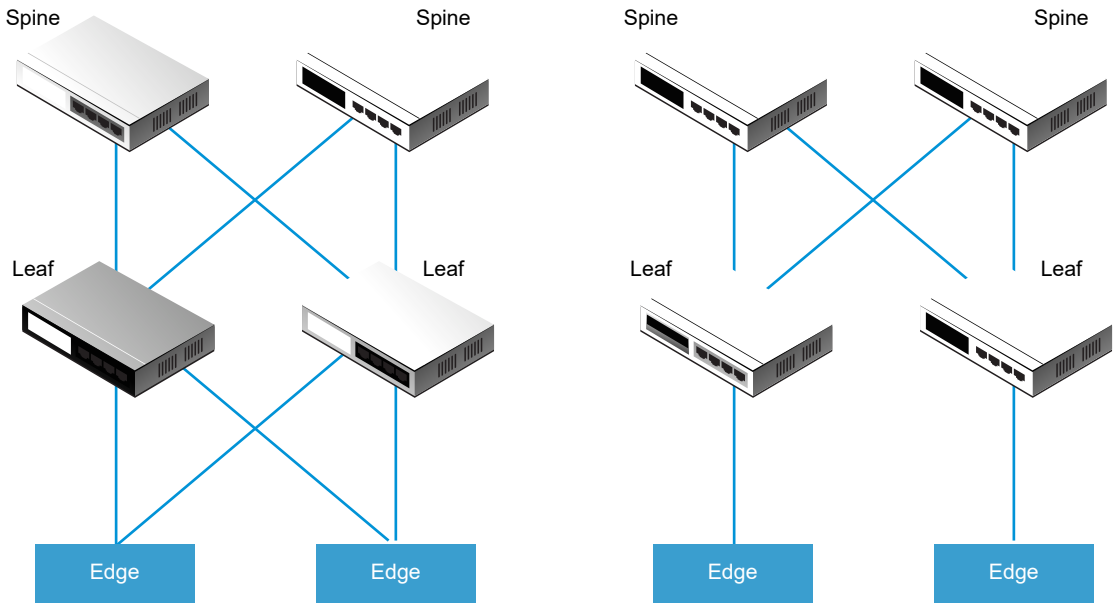
```
ping 10.10.10.10
```

```
PING 10.10.10.10 (10.10.10.10) 56(84) bytes of data.
64 bytes from 10.10.10.10: icmp_req=1 ttl=62 time=119 ms
64 bytes from 10.10.10.10: icmp_req=2 ttl=62 time=1.93 ms
^C
--- 10.10.10.10 ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 1001ms
rtt min/avg/max/mdev = 1.936/60.865/119.795/58.930 ms
```

BGP Configuration Options

To take full advantage of the tier-0 logical router, the topology must be configured with redundancy and symmetry with BGP between the tier-0 routers and the external top-of-rack peers. This design helps to ensure connectivity in the event of link and node failures.

There are two modes of configuration: active-active and active-standby. The following diagram shows two options for symmetric configuration. There are two NSX Edge nodes shown in each topology. In the case of an active-active configuration, when you create tier-0 uplink ports, you can associate each uplink port with up to eight NSX Edge transport nodes. Each NSX Edge node can have two uplinks.



For option 1, when the physical leaf-node routers are configured, they should have BGP neighborships with the NSX Edges. Route redistribution should include the same network prefixes with equal BGP metrics to all of the BGP neighbors. In the tier-0 logical router configuration, all leaf-node routers should be configured as BGP neighbors.

When you are configuring the tier-0 router's BGP neighbors, if you do not specify a local address (the source IP address), the BGP neighbor configuration is sent to all NSX Edge nodes associated with the tier-0 logical router uplinks. If you do configure a local address, the configuration goes to the NSX Edge node with the uplink owning that IP address.

In the case of option1, if the uplinks are on the same subnet on the NSX Edge nodes, it makes sense to omit the local address. If the uplinks on the NSX Edge nodes are in different subnets, the local address should be specified in the tier-0 router's BGP neighbor configuration to prevent the configuration from going to all associated NSX Edge nodes.

For option 2, ensure that the tier-0 logical router configuration includes the tier-0 services router's local IP address. The leaf-node routers are configured with only the NSX Edges that they are directly connected to as the BGP neighbor.

Configure BGP on a Tier-0 Logical Router in Manager Mode

To enable access between your VMs and the outside world, you can configure an external or internal BGP (eBGP/iBGP) connection between a tier-0 logical router and a router in your physical infrastructure.

The iBGP feature has the following capabilities and restrictions:

- Redistribution, prefix lists, and routes maps are supported.
- Route reflectors are not supported.

- BGP confederation is not supported.

When configuring BGP, you must configure a local Autonomous System (AS) number for the tier-0 logical router. For example, the following topology shows the local AS number is 64510. You must also configure the remote AS number. EBGP neighbors must be directly connected and in the same subnet as the tier-0 uplink. If they are not in the same subnet, BGP multi-hop should be used.

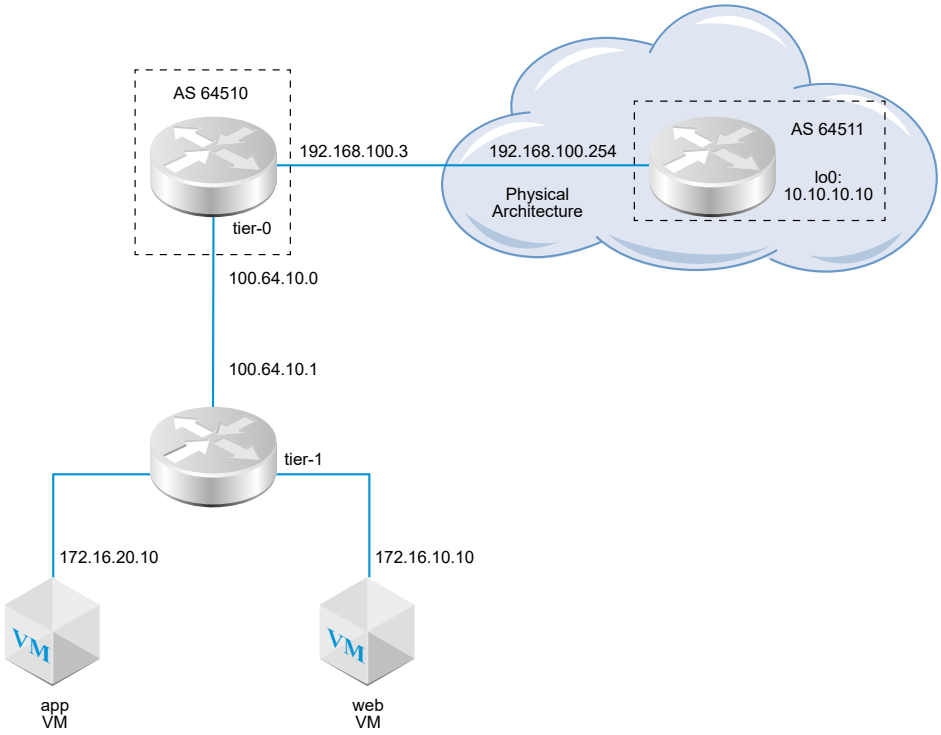
A tier-0 logical router in active-active mode supports inter-SR (service router) routing. If router #1 is unable to communicate with a northbound physical router, traffic is re-routed to router #2 in the active-active cluster. If router #2 is able to communicate with the physical router, traffic between router #1 and the physical router will not be affected.

In a topology with a tier-0 logical router in active-active mode attached to a tier-1 logical router in active-standby mode, you must enable inter-SR routing to handle asymmetric routing. You have asymmetric routing if you configure a static route on one of the SRs, or if one SR needs to reach another SR's uplink. In addition, note the following:

- In the case of a static route configured on one SR (for example, SR #1 on Edge node #1), another SR (for example, SR #2 on Edge node #2) might learn the same route from an eBGP peer and prefer the learned route to the static route on SR #1, which might be more efficient. To ensure that SR #2 uses the static route configured on SR #1, configure the tier-1 logical router in pre-emptive mode and configure Edge node #1 as the preferred node.
- If the tier-0 logical router has an uplink port on Edge node #1 and another uplink port on Edge node #2, ping traffic from tenant VMs to the uplinks works if the two uplinks are in different subnets. Ping traffic will fail if the two uplinks are in the same subnet.

Note Router ID used for forming BGP sessions on an edge node is automatically selected from the IP addresses configured on the uplinks of a tier-0 logical router. BGP sessions on an edge node can flap when router ID changes. This can happen when the IP address auto-selected for router ID is deleted or the logical router port on which this IP is assigned is deleted.

Figure 25-10. BGP Connection Topology



Caution Note the following scenarios when there are connection failures involving BGP or BFD:

- With only BGP configured, if all BGP neighbors go down, the service router's state will be down.
- With only BFD configured, if all BFD neighbors go down, the service router's state will be down.
- With BGP and BFD configured, if all BGP and BFD neighbors go down, the service router's state will be down.
- With BGP and static routes configured, if all BGP neighbors go down, the service router's state will be down.
- With only static routes configured, the service router's state will always be up unless the node is experiencing a failure or in a maintenance mode.

Prerequisites

- Verify that the tier-1 router is configured to advertise connected routes. See Configure Route Advertisement on a Tier-1 Logical Router in Manager Mode. This is not strictly a prerequisite for BGP configuration, but if you have a two-tier topology and you plan to redistribute your tier-1 networks into BGP, this step is required.
- Verify that a tier-0 router is configured. See Create a Tier-0 Logical Router in Manager Mode.

- Make sure the tier-0 logical router has learned routes from the tier-1 logical router. See Verify that a Tier-0 Router Has Learned Routes from a Tier-1 Router.
- Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Networking > Tier-0 Logical Routers**.

3 Select the tier-0 logical router.

4 Click the **Routing** tab and select **BGP** from the drop-down menu.

5 Click **Edit**.

   a Enter the local AS number.

   For example, 64510.

   b Click the **Status** toggle to enable or disable BGP.

   c Click the **ECMP** toggle to enable or disable ECMP.

   d Click the **Graceful Restart** toggle to enable or disable graceful restart.

   Graceful restart is only supported if the NSX Edge cluster associated with the tier-0 router has only one edge node.

   e If this logical router is in active-active mode, click the **Inter SR Routing** toggle to enable or disable inter-SR routing.

   f Configure route aggregation.

   g Click **Save**.

6 Click **Add** to add a BGP neighbor.

7 Enter the neighbor IP address.

   For example, 192.168.100.254.

8 Specify the maximum hop limit.

   The default is 1.

9 Enter the remote AS number.

   For example, 64511 (eBGP neighbor) or 64510 (iBGP neighbor).

10 Configure the timers (keep alive time and hold down time) and a password.

11 Click the **Local Address** tab to select a local address.

   a (Optional) Uncheck **All Uplinks** to see loopback ports as well as uplink ports.

12 Click the **Address Families** tab to add an address family.

13 Click the **BFD Configuration** tab to enable BFD.

14 Click **Save**.

What to do next

Test whether BGP is working properly. See Verify BGP Connections from a Tier-0 Service Router .

Verify BGP Connections from a Tier-0 Service Router

Use the CLI to verify from the tier-0 service router that a BGP connection to a neighbor is established.

Prerequisites

Verify that BGP is configured. See Configure BGP on a Tier-0 Logical Router in Manager Mode.

Procedure

1 Log in to the NSX Manager CLI.

2 On the NSX Edge, run the `get logical-routers` command to find the VRF number of the tier-0 service router.

```
nsx-edge-1> get logical-routers
Logical Router
UUID        : 736a80e3-23f6-5a2d-81d6-bbefb2786666
vrf         : 0
type        : TUNNEL

Logical Router
UUID        : 421a2d0d-f423-46f1-93a1-2f9e366176c8
vrf         : 5
type        : SERVICE_ROUTER_TIER0

Logical Router
UUID        : f3ce9d7d-7123-47d6-aba6-45cf1388ca7b
vrf         : 6
type        : DISTRIBUTED_ROUTER

Logical Router
UUID        : c8e64eff-02b2-4462-94ff-89f3788f1a61
vrf         : 7
type        : SERVICE_ROUTER_TIER1

Logical Router
UUID        : fb6c3f1f-599f-4421-af8a-99692dff3dd4
vrf         : 8
type        : DISTRIBUTED_ROUTER
```

3 Run the `vrf <number>` command to enter the tier-0 service router context.

```
nsx-edge-1> vrf 5
nsx-edge1(tier0_sr)>
```

4 Verify that the BGP state is `Established, up`.

`get bgp neighbor`

```
BGP neighbor: 192.168.100.254   Remote AS: 64511
BGP state: Established, up
Hold Time: 180s   Keepalive Interval: 60s
Capabilities:
         Route Refresh: advertised and received
         Address Family: IPv4 Unicast:advertised and received
         Graceful Restart: none
         Restart Remaining Time: 0
Messages: 28 received, 31 sent
Minimum time between advertisements: 30s (default)
For Address Family IPv4 Unicast:advertised and received
         Route Refresh: 0 received, 0 sent
         Prefixes: 2 received, 2 sent, 2 advertised
1 Connections established, 2 dropped
Local host: 192.168.100.3, Local port: 179
Remote host: 192.168.100.254, Remote port: 33044
```

What to do next

Check the BGP connection from the external router. See Verify North-South Connectivity and Route Redistribution on a Tier-0 Router.

Configure BFD on a Tier-0 Logical Router in Manager Mode

BFD (Bidirectional Forwarding Detection) is a protocol that can detect forwarding path failures.

Note In this release, BFD over Virtual Tunnel Interface (VTI) ports is not supported.

Caution Note the following scenarios when there are connection failures involving BGP or BFD:

- With only BGP configured, if all BGP neighbors go down, the service router's state will be down.
- With only BFD configured, if all BFD neighbors go down, the service router's state will be down.
- With BGP and BFD configured, if all BGP and BFD neighbors go down, the service router's state will be down.
- With BGP and static routes configured, if all BGP neighbors go down, the service router's state will be down.
- With only static routes configured, the service router's state will always be up unless the node is experiencing a failure or in a maintenance mode.

Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Select **Networking > Tier-0 Logical Routers**.
3. Select the tier-0 logical router.
4. Click the **Routing** tab and select **BFD** from the drop-down menu.
5. Click **Edit** to configure BFD.
6. Click the **Status** toggle to enable BFD.

   You can optionally change the global BFD properties **Receive interval**, **Transmit interval**, and **Declare dead interval**.
7. (Optional) Click **Add** under BFD Peers for Static Route Next Hops to add a BFD peer.

   Specify the peer IP address and set the admin status to **Enabled**. Optionally, you can override the global BFD properties **Receive interval**, **Transmit interval**, and **Declare dead interval**.

Enable Route Redistribution on the Tier-0 Logical Router in Manager Mode

When you enable route redistribution, the tier-0 logical router will redistribute routes into the configured destination protocol.

Prerequisites

- Verify that the tier-0 and tier-1 logical routers are connected so that you can advertise the tier-1 logical router networks to redistribute them on the tier-0 logical router. See Attach Tier-1 Router to a Tier-0 Router in Manager Mode.
- If you want to filter specific IP addresses from route redistribution, verify that route maps are configured. See Create a Route Map in Manager Mode.
- Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Select **Networking > Tier-0 Logical Routers**.
3. Select the tier-0 logical router.
4. Click the **Routing** tab and select **Route Redistribution** from the drop-down menu.
5. Click **Edit** to enable or disable route redistribution.

6 Click **Add** to add a set of route redistribution criteria.

Option	Description
Name and Description	Assign a name to the route redistribution. You can optionally provide a description. An example name, advertise-to-bgp-neighbor.
Sources	Select one or more of the following sources: ■ **T0 Connected** ■ **T0 Uplink** ■ **T0 Downlink** ■ **T0 CSP** ■ **T0 Loopback** ■ **T0 Static** ■ **T0 NAT** ■ **T0 DNS Forwarder IP** ■ **T0 IPSec Local IP** ■ **T1 Connected** ■ **T1 CSP** ■ **T1 Downlink** ■ **T1 Static** ■ **T1 LB SNAT** ■ **T1 NAT** ■ **T1 LB VIP** ■ **T1 DNS Forwarder IP**
Route Map	(Optional) Assign a route map to filter a sequence of IP addresses from route redistribution.

Verify North-South Connectivity and Route Redistribution on a Tier-0 Router

Use the CLI to verify that the BGP routes are learned. You can also check from the external router that the NSX-connected VMs are reachable.

Prerequisites

- Verify that BGP is configured. See Configure BGP on a Tier-0 Logical Router in Manager Mode.
- Verify that NSX static routes are set to be redistributed. See Enable Route Redistribution on the Tier-0 Logical Router in Manager Mode.

Procedure

1 Log in to the NSX Manager CLI.

2 View the routes learned from the external BGP neighbor.

```
nsx-edge1(tier0_sr)> get route bgp

Flags: c - connected, s - static, b - BGP, ns - nsx_static
```

```
nc - nsx_connected, rl - router_link, t0n: Tier0-NAT, t1n: Tier1-NAT


b    10.10.10.0/24        [20/0]        via 192.168.100.254
```

3 From the external router, check that BGP routes are learned and that the VMs are reachable through the NSX overlay.

a List the BGP routes.

```
user@router# run show ip route bgp
Codes: K - kernel route, C - connected, S - static, R - RIP, O - OSPF,
       I - ISIS, B - BGP, > - selected route, * - FIB route

B>* 172.16.10.0/24 [20/0] via 192.168.100.2, eth2, 00:00:48
B>* 172.16.20.0/24 [20/0] via 192.168.100.2, eth2, 00:00:48
B>* 172.16.30.0/24 [20/0] via 192.168.100.2, eth2, 00:00:48
```

b From the external router, ping the NSX-connected VMs.

`ping 172.16.10.10`

```
PING 172.16.10.10 (172.16.10.10) 56(84) bytes of data.
64 bytes from 172.16.10.10: icmp_req=1 ttl=62 time=127 ms
64 bytes from 172.16.10.10: icmp_req=2 ttl=62 time=1.96 ms
^C
--- 172.16.10.10 ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 1001ms
rtt min/avg/max/mdev = 1.966/64.793/127.620/62.827 ms
```

c Check the path through the NSX overlay.

`traceroute 172.16.10.10`

```
traceroute to 172.16.10.10 (172.16.10.10), 30 hops max, 60 byte packets
 1  192.168.100.3 (192.168.100.3)  0.640 ms  0.575 ms  0.696 ms
 2  100.91.176.1 (100.91.176.1)  0.656 ms  0.604 ms  0.578 ms
 3  172.16.10.10 (172.16.10.10)  3.397 ms  3.703 ms  3.790 ms
```

4 From the internal VMs, ping the external IP address.

`ping 10.10.10.10`

```
PING 10.10.10.10 (10.10.10.10) 56(84) bytes of data.
64 bytes from 10.10.10.10: icmp_req=1 ttl=62 time=119 ms
64 bytes from 10.10.10.10: icmp_req=2 ttl=62 time=1.93 ms
```

```
^C
--- 10.10.10.10 ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 1001ms
rtt min/avg/max/mdev = 1.936/60.865/119.795/58.930 ms
```

What to do next

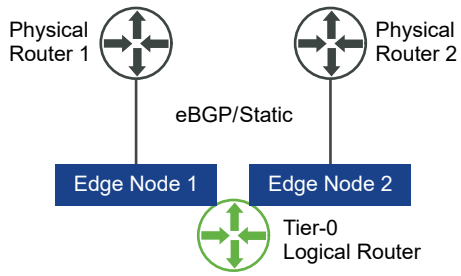
Configure additional routing functionality, such as ECMP.

Understanding ECMP Routing

Equal cost multi-path (ECMP) routing protocol increases the north and south communication bandwidth by adding an uplink to the tier-0 logical router and configure it for each Edge node in an NSX Edge cluster. The ECMP routing paths are used to load balance traffic and provide fault tolerance for failed paths.

The tier-0 logical router must be in active-active mode for ECMP to be available. A maximum of eight ECMP paths are supported. The implementation of ECMP on NSX Edge is based on the 5-tuple of the protocol number, source address, destination address, source port, and destination port. The algorithm used to distribute the data among the ECMP paths is not round robin. Therefore, some paths might carry more traffic than others. Note that if the protocol is IPv6 and the IPv6 header has more than one extension header, ECMP will be based only on the source and destination addresses.

Figure 25-11. ECMP Routing Topology



For example, the topology above shows a single tier-0 logical router in active-active mode running on a 2-node NSX Edge cluster. Two uplink ports are configured, one on each Edge node.

Add an Uplink Port for the Second Edge Node for ECMP in Manager Mode

Before you enable ECMP, you must configure an uplink to connect the tier-0 logical router to the VLAN logical switch.

Prerequisites

- Verify that a transport zone and two transport nodes are configured. See the *NSX Installation Guide*.
- Verify that two Edge nodes and an Edge cluster are configured. See the *NSX Installation Guide*.

- Verify that a VLAN logical switch for uplink is available. See Create a VLAN Logical Switch for the NSX Edge Uplink in Manager Mode.
- Verify that a tier-0 logical router is configured. See Create a Tier-0 Logical Router in Manager Mode.
- Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Networking > Tier-0 Logical Routers**.

3 Select the tier-0 logical router.

4 Click the **Configuration** tab to add a router port.

5 Click **Add**.

6 Complete the router port details.

Option	Description
Name	Assign a name for the router port.
Description	Provide additional description that shows that the port is for ECMP configuration.
Type	Accept the default type **Uplink**.
MTU	If you leave this field empty, the default is 1500.
Transport Node	Assign the Edge transport node from the drop-down menu.
URPF Mode	uRPF (unicast Reverse Path Forwarding) is enabled by default on external, internal and service interfaces. From a security standpoint, it is a best practice to keep uRPF enabled on these interfaces. uRPF is also recommended in architectures that leverage ECMP. It is possible to disable uRPF in complex routing architectures where asymmetric routing exists.
Logical Switch	Assign the VLAN logical switch from the drop-down menu.
Logical Switch Port	Assign a new switch port name. You can also use an existing switch port.
IP Address/Mask	Enter an IP address that is in the same subnet as the connected port on the ToR switch.

7 Click **Save**.

Results

A new uplink port is added to the tier-0 router and the VLAN logical switch. The tier-0 logical router is configured on both of the edge nodes.

What to do next

Create a BGP connection for the second neighbor and enable the ECMP routing. See Add a Second BGP Neighbor and Enable ECMP Routing in Manager Mode.

Add a Second BGP Neighbor and Enable ECMP Routing in Manager Mode

Before you enable ECMP routing, you must add a BGP neighbor and configure it with the newly added uplink information.

Prerequisites

- Verify that the second edge node has an uplink port configured. See Add an Uplink Port for the Second Edge Node for ECMP in Manager Mode .
- Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Select **Networking > Tier-0 Logical Routers**.
3. Select the tier-0 logical router.
4. Click the **Routing** tab and select **BGP** from the drop-down menu.
5. Click **Add** under the Neighbors section to add a BGP neighbor.
6. Enter the neighbor IP address.

   For example, 192.168.200.254.
7. (Optional) Specify the maximum hop limit.

   The default is 1.
8. Enter the remote AS number.

   For example, 64511.
9. (Optional) Click the **Local Address** tab to select a local address.

   a (Optional) Uncheck **All Uplinks** to see loopback ports as well as uplink ports.
10. (Optional) Click the **Address Families** tab to add an address family.
11. (Optional) Click the **BFD Configuration** tab to enable BFD.
12. Click **Save**.

    The newly added BGP neighbor appears.
13. Click **Edit** next to the BGP Configuration section.

14 Click the **ECMP** toggle button to enable ECMP.

The Status button must be appear as Enabled.

15 Click **Save**.

Results

Multiple ECMP routing paths connect the VMs attached to logical switches and the two Edge nodes in the Edge cluster.

What to do next

Test whether the ECMP routing connections are working properly. See Verify North-Bound ECMP Routing Connectivity on a Tier-0 Router.

Verify North-Bound ECMP Routing Connectivity on a Tier-0 Router

Procedure

1 Log in to NSX Edge CLI with the admin credentials.

2 Get the service router information.

```
get logical-routers
```

3 Type the VRF for the tier-0 service router.

```
vrf <vrf-no>
```

4 Identify the BGP peers. See if SR has at least two equal cost BGP peers.

```
get bgp neighbor summary
```

5 Identify the BGP prefixes.

The `get bgp` command will list all the prefixes in the BGP table.

The `get bgp neighbor <neighbor-IP> route` will list prefix specific to the BGP peer.

Note Ensure that BGP prefixes, which are expected to have ECMP towards the BGP peers, are learned from each of the peers.

6 Identify a specific route to confirm that it has two next hops.

```
get route
```

7 Identify the forwarding table on the Tier0 to confirm that each prefix learned with ECMP through BGP is installed with two next hops.

```
get forwarding
```

Create an IP Prefix List in Manager Mode

An IP prefix list contains single or multiple IP addresses that are assigned access permissions for route advertisement. The IP addresses in this list are processed sequentially. IP prefix lists are referenced through BGP neighbor filters or route maps with in or out direction.

For example, you can add the IP address 192.168.100.3/27 to the IP prefix list and deny the route from being redistributed to the northbound router. You can also append an IP address with less-than-or-equal-to (le) and greater-than-or-equal-to (ge) modifiers to grant or limit route redistribution. For example, 192.168.100.3/27 ge 24 le 30 modifiers match subnet masks greater than or equal to 24-bits and less than or equal to 30-bits in length.

Note The default action for a route is **Deny**. When you create a prefix list to deny or permit specific routes, be sure to create an IP prefix with no specific network address (select **Any** from the dropdown list) and the **Permit** action if you want to permit all other routes.

Prerequisites

- Verify that you have a tier-0 logical router configured. See Create a Tier-0 Logical Router in Manager Mode.
- Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Select **Networking > Tier-0 Logical Routers**.
3. Select the tier-0 logical router.
4. Click the **Routing** tab and select **IP Prefix Lists** from the drop-down menu.
5. Click **Add**.
6. Enter a name for the IP prefix list.
7. Click **Add** to specify a prefix.
   a. Enter an IP address in CIDR format.

      For example, 192.168.100.3/27.

   b. Select **Deny** or **Permit** from the drop-down menu.

   c. (Optional) Set a range of IP address numbers in the **le** or **ge** modifiers.

      For example, set **le** to 30 and **ge** to 24.
8. Repeat the previous step to specify additional prefixes.
9. Click **Add** at the bottom of the window.

Create a Community List in Manager Mode

You can create BGP community lists so that you can configure route maps based on community lists.

Prerequisites

- Verify that you have a tier-0 logical router configured. See Create a Tier-0 Logical Router in Manager Mode.
- Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Select **Networking > Tier-0 Logical Routers**.
3. Select the tier-0 logical router.
4. Click the **Routing** tab and select **Community Lists** from the drop-down menu.
5. Click **Add**.
6. Enter a name for the community list.
7. Specify a community using the aa:nn format, for example, 300:500, and press Enter. Repeat to add additional communities.

   In addition, you can click the dropdown arrow and select one or more of the following:

   - NO_EXPORT_SUBCONFED - Do not advertise to EBGP peers.
   - NO_ADVERTISE - Do not advertise to any peer.
   - NO_EXPORT - Do not advertise outside BGP confederation
8. Click **Add**.

Create a Route Map in Manager Mode

A route map consists of a sequence of IP prefix lists, BGP path attributes, and an associated action. The router scans the sequence for an IP address match. If there is a match, the router performs the action and scans no further.

Route maps can be referenced at the BGP neighbor level and route redistribution. When IP prefix lists are referenced in route maps and the route map action of permitting or denying is applied, the action specified in the route map sequence overrides the specification within the IP prefix list.

Prerequisites

- Verify that an IP prefix list is configured. See Create an IP Prefix List in Manager Mode.
- Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1. With admin privileges, log in to NSX Manager.

2 Select **Networking > Tier-0 Logical Routers**.

3 Select the tier-0 logical router.

4 Select **Routing > Route Maps**.

5 Click **Add**.

6 Enter a name and an optional description for the route map.

7 Click **Add** to add an entry in the route map.

8 Edit the column **Match IP Prefix List/Community List** to select either IP prefix lists, or community lists, but not both.

9 (Optional) Set BGP attributes.

BGP Attribute	Description
AS-path Prepend	Prepend a path with one or more AS (autonomous system) numbers to make the path longer and therefore less preferred.
MED	Multi-Exit Discriminator indicates to an external peer a preferred path to an AS.
Weight	Set a weight to influence path selection. The range is 0 - 65535.
Community	Specify a community using the aa:nn format, for example, 300:500. Or use the drop-down menu to select one of the following: ■ NO_EXPORT_SUBCONFED - Do not advertise to EBGP peers. ■ NO_ADVERTISE - Do not advertise to any peer. ■ NO_EXPORT - Do not advertise outside BGP confederation

10 In the Action column, select **Permit** or **Deny**.

You can permit or deny IP addresses in the IP prefix lists from advertising their addresses.

11 Click **Save**.

Configure Forwarding Up Timer in Manager Mode

You can configure forwarding up timer for a tier-0 logical router.

Forwarding up timer defines the time in seconds that the router must wait before sending the up notification after the first BGP session is established. This timer (previously known as forwarding delay) minimizes downtime in case of fail-overs for active-active or active-standby configurations of logical routers on NSX Edge that use dynamic routing (BGP). It should be set to the number of seconds an external router (TOR) takes to advertise all the routes to this router after the first BGP/BFD session. The timer value should be directly proportional to the number of northbound dynamic routes that the router must learn. This timer should be set to 0 on single edge node setups.

Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Networking > Tier-0 Logical Routers**.

3 Select the tier-0 logical router.

4 Select **Routing > Global Configuration**

5 Click **Edit**.

6 Enter a value for the forwarding up timer.

7 Click **Save**.

NAT in Manager Mode

You can configure Network Address Translation (NAT) in **Manager** mode.

Note If you use **Manager** mode to modify objects created in the **Policy** mode, some settings might not be configurable. These read-only settings have this icon next to them: ⊖. See Chapter 1 NSX Manager for more information.

Network Address Translation

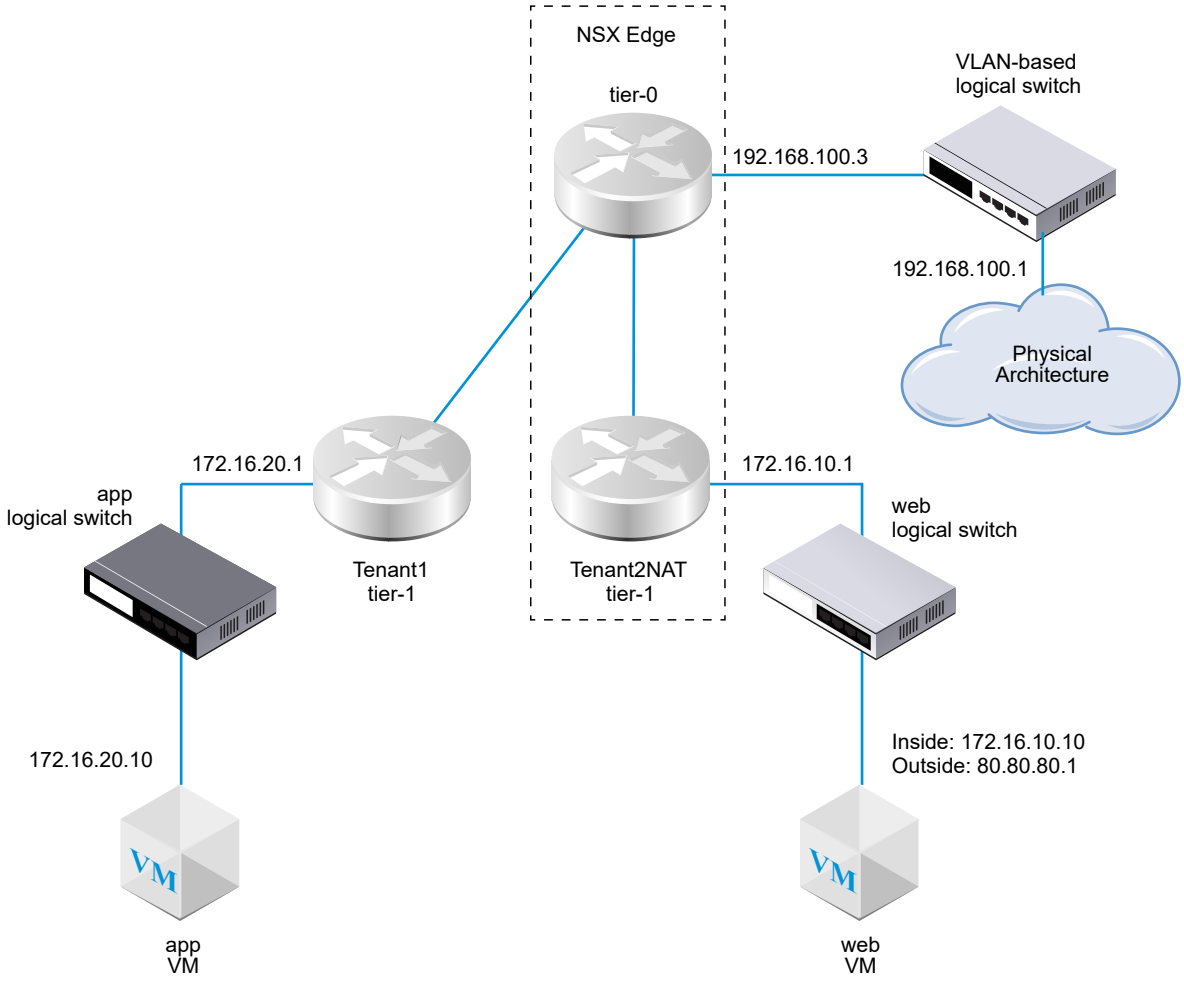
Network address translation (NAT) in NSX can be configured on tier-0 and tier-1 logical routers.

For example, the following diagram shows two tier-1 logical routers with NAT configured on Tenant2NAT. The web VM is simply configured to use 172.16.10.10 as its IP address and 172.16.10.1 as its default gateway.

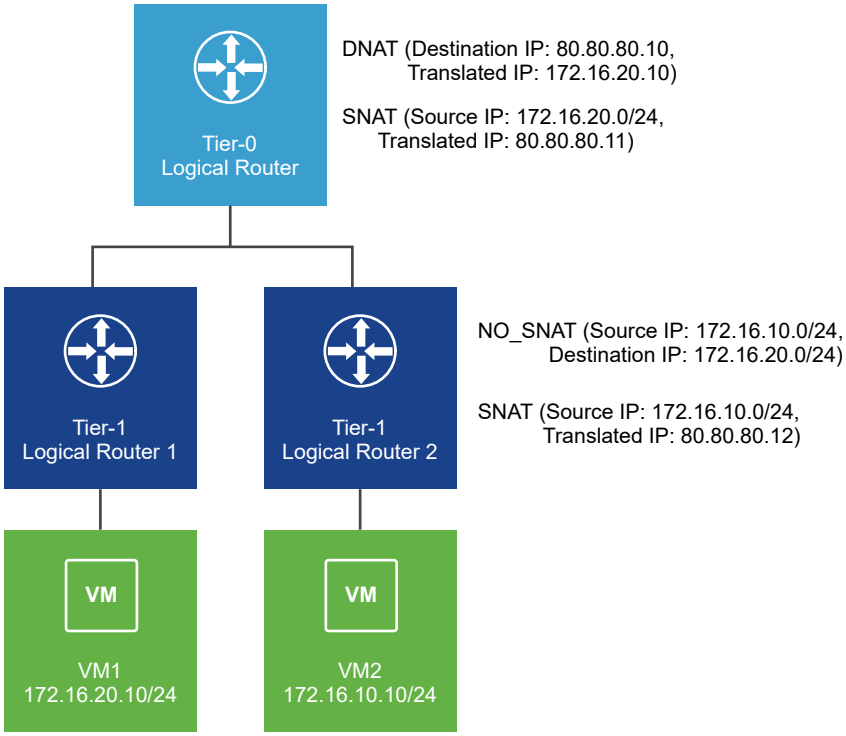
NAT is enforced at the uplink of the Tenant2NAT logical router on its connection to the tier-0 logical router.

To enable NAT configuration, Tenant2NAT must have a service component on an NSX Edge cluster. Thus, Tenant2NAT is shown inside the NSX Edge. For comparison, Tenant1 can be outside of the NSX Edge because it is not using any Edge services.

Figure 25-12. NAT Topology



Note: In the following scenario, NAT hairpinning is not supported. The tier-0 logical router has DNAT and SNAT configured. Tier-1 Logical Router 2 has NO_SNAT and SNAT configured. VM2 will not be able to access VM1 using VM1's external address 80.80.80.10.



The following sections describe how to create NAT rules using the manager UI. You can alco make an API call (`POST /api/v1/logical-routers/<logical-router-id>/nat/rules?action=create_multiple`) to create multiple NAT rules at the same time. For more information, see the *NSX API Guide*.

Tier-1 NAT

A tier-1 logical router supports source NAT (SNAT), destination NAT (DNAT) and reflexive NAT.

Configure Source NAT on a Tier-1 Router in Manager Mode

Source NAT (SNAT) changes the source address in the IP header of a packet. It can also change the source port in the TCP/UDP headers. The typical usage is to change a private (rfc1918) address/port into a public address/port for packets leaving your network.

You can create a rule to either enable or disable source NAT.

In this example, as packets are received from the web VM, the Tenant2NAT tier-1 router changes the source IP address of the packets from 172.16.10.10 to 80.80.80.1. Having a public source IP address enables destinations outside of the private network to route back to the original source.

Prerequisites

- The tier-0 router must have an uplink connected to a VLAN-based logical switch. See Connect a Tier-0 Logical Router to a VLAN Logical Switch for the NSX Edge Uplink in Manager Mode.

- The tier-0 router must have routing (static or BGP) and route redistribution configured on its uplink to the physical architecture. See Configure a Static Route in Manager Mode, Configure BGP on a Tier-0 Logical Router in Manager Mode, and Enable Route Redistribution on the Tier-0 Logical Router in Manager Mode.
- The tier-1 routers must each have an uplink to a tier-0 router configured. Tenant2NAT must be backed by an NSX Edge cluster. See Attach Tier-1 Router to a Tier-0 Router in Manager Mode.
- The tier-1 routers must have downlink ports and route advertisement configured. See Add a Downlink Port on a Tier-1 Logical Router in Manager Mode and Configure Route Advertisement on a Tier-1 Logical Router in Manager Mode.
- The VMs must be attached to the correct logical switches.
- Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Select **Networking > Tier-0 Logical Routers**.
3. Click a tier-1 logical router on which you want to configure NAT.
4. Select **Services > NAT**.
5. Click **ADD**.
6. Specify a priority value.

   A lower value means a higher precedence for this rule.
7. For **Action**, select **SNAT** to enable source NAT, or **NO_SNAT** to disable source NAT.
8. Select the protocol type.

   By default, **Any Protocol** is selected.
9. (Optional) For **Source IP**, specify an IP address or an IP address range in CIDR format.

   If you leave this field blank, all sources on router's downlink ports are translated. In this example, the source IP address is 172.16.10.10.
10. (Optional) For **Destination IP**, specify an IP address or an IP address range in CIDR format.

    If you leave this field blank, the NAT applies to all destinations outside of the local subnet.
11. If **Action** is **SNAT**, for **Translated IP**, specify an IP address or an IP address range in CIDR format.

    In this example, the translated IP address is 80.80.80.1.
12. (Optional) For **Applied To**, select a router port.

13 (Optional) Set the status of the rule.

The rule is enabled by default.

14 (Optional) Change the logging status.

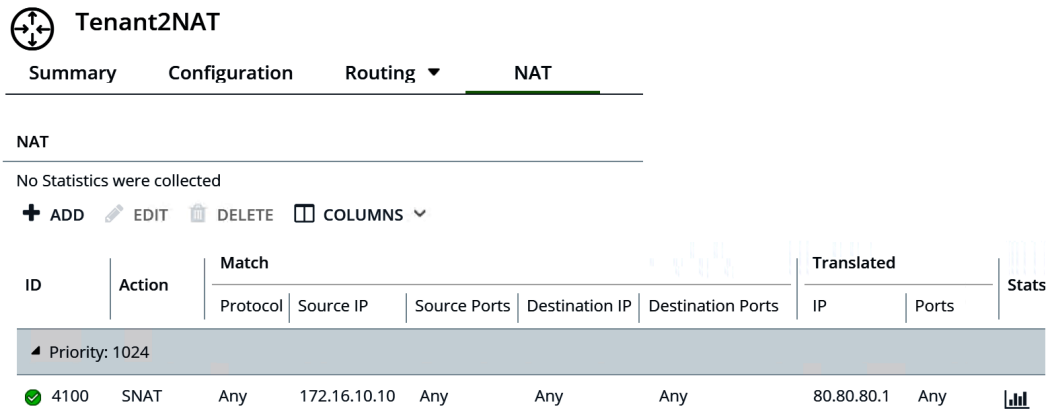
Logging is disabled by default.

15 (Optional) Change the firewall bypass setting.

The setting is enabled by default.

Results

The new rule is listed under NAT. For example:



What to do next

Configure the tier-1 router to advertise NAT routes.

To advertise the NAT routes upstream from the tier-0 router to the physical architecture, configure the tier-0 router to advertise tier-1 NAT routes.

Configure Destination NAT on a Tier-1 Router in Manager Mode

Destination NAT changes the destination address in IP header of a packet. It can also change the destination port in the TCP/UDP headers. The typical usage of this is to redirect incoming packets with a destination of a public address/port to a private IP address/port inside your network.

You can create a rule to either enable or disable destination NAT.

In this example, as packets are received from the app VM, the Tenant2NAT tier-1 router changes the destination IP address of the packets from 172.16.10.10 to 80.80.80.1. Having a public destination IP address enables a destination inside a private network to be contacted from outside of the private network.

Prerequisites

- The tier-0 router must have an uplink connected to a VLAN-based logical switch. See Connect a Tier-0 Logical Router to a VLAN Logical Switch for the NSX Edge Uplink in Manager Mode.

- The tier-0 router must have routing (static or BGP) and route redistribution configured on its uplink to the physical architecture. See Configure a Static Route in Manager Mode, Configure BGP on a Tier-0 Logical Router in Manager Mode, and Enable Route Redistribution on the Tier-0 Logical Router in Manager Mode.
- The tier-1 routers must each have an uplink to a tier-0 router configured. Tenant2NAT must be backed by an NSX Edge cluster. See Attach Tier-1 Router to a Tier-0 Router in Manager Mode.
- The tier-1 routers must have downlink ports and route advertisement configured. See Add a Downlink Port on a Tier-1 Logical Router in Manager Mode and Configure Route Advertisement on a Tier-1 Logical Router in Manager Mode.
- The VMs must be attached to the correct logical switches.
- Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Select **Networking > Tier-1 Logical Routers**.
3. Click a tier-1 logical router on which you want to configure NAT.
4. Select **Services > NAT**.
5. Click **ADD**.
6. Specify a priority value.

   A lower value means a higher precedence for this rule.
7. For **Action**, select **DNAT** to enable destination NAT, or **NO_DNAT** to disable destination NAT.
8. Select the protocol type.

   By default, **Any Protocol** is selected.
9. (Optional) For **Source IP**, specify an IP address or an IP address range in CIDR format.

   If you leave Source IP blank, the NAT applies to all sources outside of the local subnet.
10. For **Destination IP**, specify an IP address or a comma-separated IP address list.

    In this example, the destination IP address is 80.80.80.1.
11. If **Action** is **DNAT**, for **Translated IP**, specify an IP address or an IP address range in CIDR format.

    In this example, the inside/translated IP address is 172.16.10.10.
12. (Optional) If **Action** is **DNAT**, for **Translated Ports**, specify the translated ports.

13 (Optional) For **Applied To**, select a router port.

14 (Optional) Set the status of the rule.

The rule is enabled by default.

15 (Optional) Change the logging status.

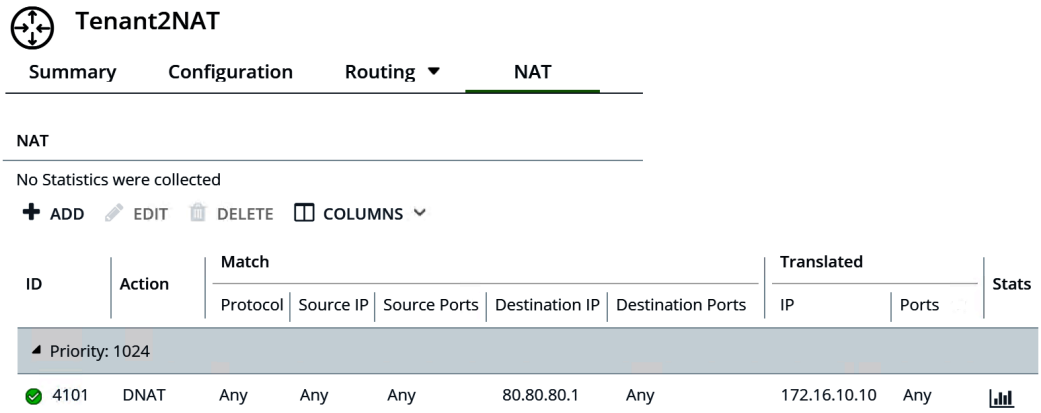
Logging is disabled by default.

16 (Optional) Change the firewall bypass setting.

The setting is enabled by default.

Results

The new rule is listed under NAT. For example:



What to do next

Configure the tier-1 router to advertise NAT routes.

To advertise the NAT routes upstream from the tier-0 router to the physical architecture, configure the tier-0 router to advertise tier-1 NAT routes.

Advertise Tier-1 NAT Routes to the Upstream Tier-0 Router in Manager Mode

Advertising tier-1 NAT routes enables the upstream tier-0 router to learn about these routes.

Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Networking > Tier-1 Logical Routers**.

3 Click a tier-1 logical router on which you have configured NAT.

4 From the tier-1 router, select **Routing > Route Advertisement**.

5 Click **Edit** to edit the route advertisement configuration.

You can toggle the following switches:

- **Status**
- **Advertise All NSX Connected Routes**
- **Advertise All NAT Routes**
- **Advertise All Static Routes**
- **Advertise All LB VIP Routes**
- **Advertise All LB SNAT IP Routes**
- **Advertise All DNS Forwarder Routes**

6 Click **Save**.

What to do next

Advertise tier-1 NAT routes from the tier-0 router to the upstream physical architecture.

Advertise Tier-1 NAT Routes to the Physical Architecture in Manager Mode

Advertising tier-1 NAT routes from the tier-0 router enables the upstream physical architecture to learn about these routes.

Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Networking > Tier-0 Logical Routers**.

3 Click a tier-0 logical router that is connected to a tier-1 router on which you have configured NAT.

4 From the tier-0 router, select **Routing > Route Redistribution**.

5 Click **Edit** to enable or disable route redistribution.

6 Click **Add** to add a set of route redistribution criteria.

Option	Description
Name and Description	Assign a name to the route redistribution. You can optionally provide a description. An example name, advertise-to-bgp-neighbor.
Sources	Select one or more of the following sources: ■ **T0 Connected** ■ **T0 Uplink** ■ **T0 Downlink** ■ **T0 CSP** ■ **T0 Loopback** ■ **T0 Static** ■ **T0 NAT** ■ **T0 DNS Forwarder IP** ■ **T0 IPSec Local IP** ■ **T1 Connected** ■ **T1 CSP** ■ **T1 Downlink** ■ **T1 Static** ■ **T1 LB SNAT** ■ **T1 NAT** ■ **T1 LB VIP** ■ **T1 DNS Forwarder IP**
Route Map	(Optional) Assign a route map to filter a sequence of IP addresses from route redistribution.

Verify Tier-1 NAT

Verify that SNAT and DNAT rules are working correctly.

Procedure

1 Log in the NSX Edge.

2 Run `get logical-routers` to determine the VRF number for the tier-0 services router.

3 Enter the tier-0 services router context by running the `vrf <number>` command.

4 Run the `get route` command and make sure that the tier-1 NAT address appears.

```
nsx-edge(tier0_sr)> get route

Flags: c - connected, s - static, b - BGP, ns - nsx_static
nc - nsx_connected, rl - router_link, t0n: Tier0-NAT, t1n: Tier1-NAT
```

```
Total number of routes: 8

t1n  80.80.80.1/32        [3/3]         via 169.0.0.1
...
```

5 If your Web VM is set up to serve Web pages, make sure you can open a Web page at http://80.80.80.1.

6 Make sure that the tier-0 router's upstream neighbor in the physical architecture can ping 80.80.80.1.

7 While the ping is still running, check the stats column for the DNAT rule.

There should be one active session.

Tier-0 NAT

A tier-0 logical router in active-standby mode supports source NAT (SNAT), destination NAT (DNAT) and reflexive NAT. A tier-0 logical router in active-active mode supports reflexive NAT only.

Configure Source and Destination NAT on a Tier-0 Logical Router in Manager Mode

You can configure source and destination NAT on a tier-0 logical router that is running in active-standby mode.

You can also disable SNAT or DNAT for an IP address or a range of addresses. If multiple NAT rules apply to an address, the rule with the highest priority is applied.

SNAT configured on a tier-0 logical router's uplink will process traffic from a tier-1 logical router as well as from another uplink on the tier-0 logical router.

Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Networking > Tier-0 Logical Routers**.

3 Click a tier-0 logical router.

4 Select **Services > NAT**.

5 Click **ADD** to add a NAT rule.

6 Specify a priority value.

A lower value means a higher priority.

7 For **Action**, select **SNAT**, **DNAT**, **Reflexive**, **NO_SNAT**, or **NO_DNAT**.

8 Select the protocol type.

   By default, **Any Protocol** is selected.

9 (Required) For **Source IP**, specify an IP address or an IP address range in CIDR format.

   If you leave this field blank, this NAT rule applies to all sources outside of the local subnet.

10 For **Destination IP**, specify an IP address or an IP address range in CIDR format.

11 For **Translated IP**, specify an IP address or an IP address range in CIDR format.

12 (Optional) If **Action** is **DNAT**, for **Translated Ports**, specify the translated ports.

13 (Optional) For **Applied To**, select a router port.

14 (Optional) Set the status of the rule.

   The rule is enabled by default.

15 (Optional) Change the logging status.

   Logging is disabled by default.

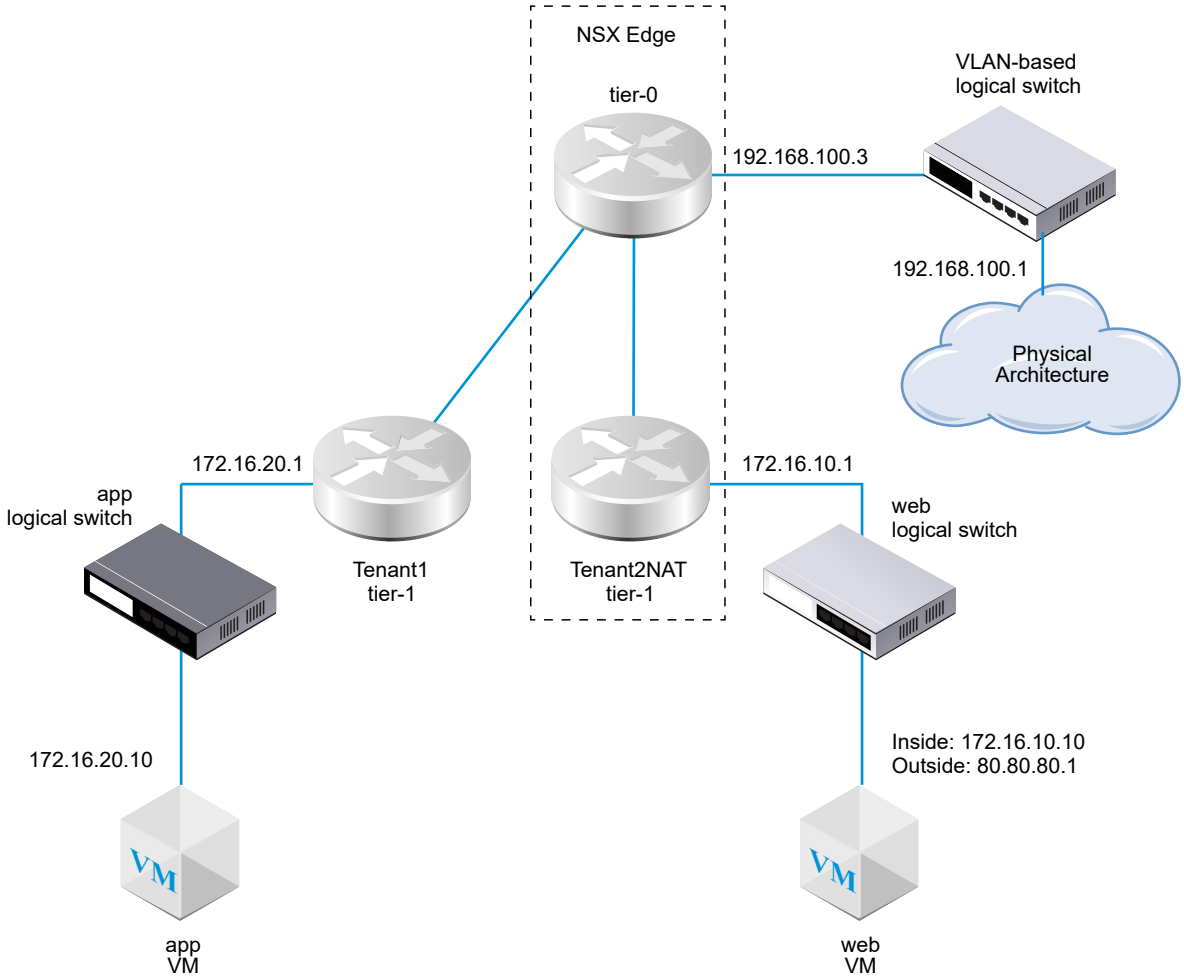
16 (Optional) Change the firewall bypass setting.

   The setting is enabled by default.

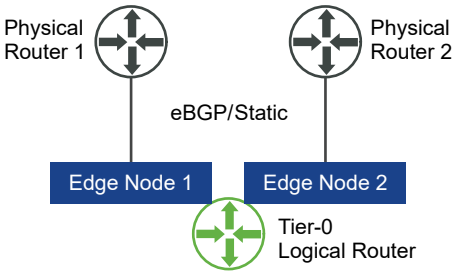
Reflexive NAT

When a tier-0 logical router is running in active-active mode, you cannot configure stateful NAT where asymmetrical paths might cause issues. For active-active routers, you can configure reflexive NAT (sometimes called stateless NAT).

In this example, as packets are received from the web VM, the Tenant2NAT tier-1 router changes the source IP address of the packets from 172.16.10.10 to 80.80.80.1. Having a public source IP address enables destinations outside of the private network to route back to the original source.



When there are two active-active tier-0 routers involved, as shown below, reflexive NAT must be configured.



Configure Reflexive NAT on a Tier-0 or Tier-1 Logical Router in Manager Mode

When a tier-0 or tier-1 logical router is running in active-active mode, you cannot configure stateful NAT where asymmetrical paths might cause issues. For active-active routers, you can use reflexive NAT, which is sometimes called stateless NAT.

For reflexive NAT, you can configure a single source address to be translated, or a range of addresses. If you configure a range of source addresses, you must also configure a range of translated addresses. The size of the two ranges must be the same. The address translation will be deterministic, meaning that the first address in the source address range will be translated to the first address in the translated address range, the second address in the source range will be translated to the second address in the translated range, and so on.

Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Locate the logical router you want to modify in **Networking > Tier-0 Logical Routers** or **Networking > Tier-1 Logical Routers**.
3. Click the tier-0 or tier-1 logical router on which you want to configure reflexive NAT.
4. Select **Services > NAT**.
5. Click **ADD**.
6. Specify a priority value.

   A lower value means a higher precedence for this rule.
7. For **Action**, select **Reflexive**.
8. For **Source IP**, specify an IP address or an IP address range in CIDR format.
9. For **Translated IP**, specify an IP address or an IP address range in CIDR format.
10. (Optional) Set the status of the rule.

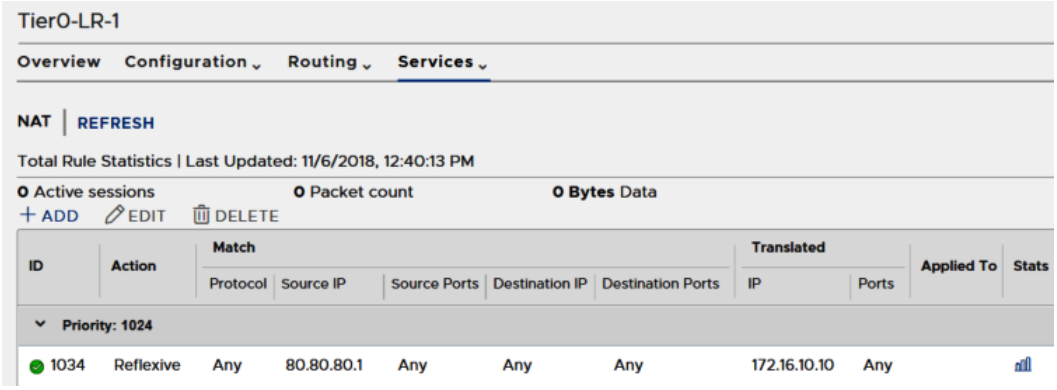
    The rule is enabled by default.
11. (Optional) Change the logging status.

    Logging is disabled by default.
12. (Optional) Change the firewall bypass setting.

    The setting is enabled by default.

Results

The new rule is listed under NAT. For example:



Grouping Objects in Manager Mode

You can create IP sets, IP pools, MAC sets, NSGroups, and NSServices in **Manager** mode. You can also manage tags for VMs.

For Policy Groups containing IPs, MAC addresses, and Identity Groups the listing API will **NOT** display the 'members' attribute. This applies to Groups containing a combination of static members also. For example, a Policy Group containing IP and VMs, will not display the the members attribute.

For Policy Groups not containing IPs, MAC addresses, or Identity Groups, the member attribute will be displayed in the NSGroup response. However new members and criteria introduced in NSX (such as DVPort and DVPG) will not be included in the MP group definition. Users can view the definition in Policy.

Note If you use **Manager** mode to modify objects created in the **Policy** mode, some settings might not be configurable. These read-only settings have this icon next to them: ⊜. See Chapter 1 NSX Manager for more information.

Create an IP Set in Manager Mode

An IP set is a group of IP addresses that can be used as sources and destinations in firewall rules.

An IP set can contain a combination of individual IP addresses, IP ranges, and subnets. You can specify IPv4 or IPv6 addresses, or both. An IP set can be a member of NSGroups.

Note Any IP set created by this method will not be visible in Policy mode. In Policy mode, we can create a group and add members as IP addresses, ranges, network addresses, or MAC addresses by navigating to **Inventory > Groups > Set Members** and specifying IP or MAC addresses.

Note IPv4 addresses and IPv6 addresses are supported for source or destination ranges for firewall rules.

Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Inventory > Groups > IP Sets > Add**.

3 Enter a name.

4 (Optional) Enter a description.

5 In **Members**, enter individual IP addresses, IP ranges, and subnets in a comma separated list.

6 Click **Save**.

Create an IP Pool in Manager Mode

You can use an IP Pool to allocate IP addresses or subnets when you create L3 subnets.

Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Networking > IP Management > IP Address Pools**.

3 Enter a name for the new IP pool.

4 (Optional) Enter a description.

5 Click **Add**.

6 Click the IP Ranges cell and enter IP Ranges.

Mouse over the upper right corner of any cell and click the pencil icon to edit it.

7 (Optional) Enter a Gateway.

8 Enter a CIDR IP address with suffix.

9 (Optional) Enter DNS Servers.

10 (Optional) Enter a DNS Suffix.

11 Click **Save**.

Create a MAC Set in Manager Mode

A MAC Set is a group of MAC addresses that you can use as sources and destinations in layer 2 firewall rules and as a member of an NS Group.

Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Select **Inventory > Groups > MAC Sets > Add**.
3. Enter a name.
4. (Optional) Enter a description.
5. Enter the MAC addresses in a comma-separated list.
6. Click **ADD**.

Create an NSGroup in Manager Mode

NSGroups can be configured to contain a combination of IP sets, MAC sets, logical ports, logical switches, and other NSGroups. You can specify NSGroups with Logical Switches, Logical ports and VMs as sources and destinations, and in the `Applied To` field of a firewall rule. NSGroups with IPset and MACSet will be ignored in a distributed firewall `Applied To` field.

NSX Cloud Note If using NSX Cloud, see NSX Features Supported with NSX Cloud for a list of auto-generated logical entities, supported features, and configurations required for NSX Cloud.

An NSGroup has the following characteristics:

- An NSGroup has direct members and effective members. Effective members include members that you specify using membership criteria, as well as all the direct and effective members that belong to this NSGroup's members. For example, assuming NSGroup-1 has direct member LogicalSwitch-1. You add NSGroup-2 and specify NSGroup-1 and LogicalSwitch-2 as members. Now NSGroup-2 has direct members NSGroup-1 and LogicalSwitch-2, and an effective member, LogicalSwitch-1. Next, you add NSGroup-3 and specify NSGroup-2 as a member. NSGroup-3 now has direct member NSGroup-2 and effective members LogicalSwitch-1 and LogicalSwitch-2. From the main groups table, clicking on a group and selecting **Related > NSGroups** would show NSGroup-1, NSGroup-2, and NSGroup-3 because all three have LogicalSwitch-1 as a member, either directly or indirectly.
- An NSGroup can have a maximum of 500 direct members.

- The recommended limit for the number of effective members in an NSGroup is 5000. The NSX Manager check the NSGroups regarding the limit twice a day, at 7 AM and 7 PM. Exceeding this limit does not affect any functionality but might have a negative impact on performance.
  - When the number of effective members for an NSGroup exceeds 80% of 5000, the warning message `NSGroup xyz is about to exceed the maximum member limit. Total number in NSGroup is ...` appears in the log file. When the number exceeds 5000, the warning message `NSGroup xyz has reached the maximum numbers limit. Total number in NSGroup = ...` appears.
  - When the number of translated VIFs/IPs/MACs in an NSGroup exceeds 5000, the warning message `Container xyz has reached the maximum IP/MAC/VIF translations limit. Current translations count in Container - IPs:..., MACs:..., VIFs:...` appears in the log file.
- The maximum supported number of VMs is 10,000.
- You can create a maximum of 10,000 NSGroups.
- Edge_NSGroup is a policy owned group (system group) which is available on a local manager and is visible on the UI. This group is not available on a global manager. However, a migrated global manager setup contains stale Edge_NSGroup and UI displays the same, but the group holds no significance on a global manager.

For all the objects that you can add to an NSGroup as members, you can navigate to the screen for any of the objects and select **Related > NSGroups**.

Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Select **Inventory > Groups > Add**.
3. Enter a name for the NSGroup.
4. (Optional) Enter a description.
5. (Optional) Click **Membership Criteria**.

   For each criterion, you can specify up to five rules, which are combined with the logical AND operator. The available member criterion can apply to the following:

   - **Logical Port** - can specify a tag and optional scope.
   - **Logical Switch** - can specify a tag and optional scope.

- **Virtual Machine** - can specify a name, tag, computer OS name, or computer name that equals, contains, starts with, ends with, or doesn't equal a particular string.
- **Transport Node** - can specify a node type that equals an edge node or a host node.
- **IP Set** - can specify a tag and optional scope.

6 (Optional) Click **Members** to select members.

The available member types are:

- **AD Group** - NSGroups with ADGroups can only be used in the extended_source field of a distributed firewall rule, and must be the only members in the group. For example, there cannot be an NSGroup with both ADGroup and IPSet together as members.
- **IP Set** - can include both IPv4 an IPv6 addresses.
- **Logical Port** - can include both IPv4 and IPv6 addresses.
- **Logical Switch** - can include both IPv4 and IPv6 addresses.
- **MAC Set**
- **NSGroup**
- **Transport Node**
- **VIF**
- **Virtual Machine**

7 Click **ADD**.

The group is added to the table of groups. Click a group name to display an overview and edit group information including membership criteria, members, applications, and related groups. Scroll to the bottom of the **Overview** tab to add and delete tags. See Add Tags to an Object for more information. Selecting **Related> NSGroups** displays all the NSGroups that have the selected NSGroup as a member.

Configuring Services and Service Groups

You can configure an NSService and specify parameters for matching network traffic such as a port and protocol pairing. You can also use an NSService to allow or block certain types of traffic in firewall rules.

An NSService can be of the following types:

- Ether
- IP
- IGMP
- ICMP
- ALG
- L4 Port Set

An L4 Port Set supports the identification of source ports and destination ports. You can specify individual ports or a range of ports, up to a maximum of 15 ports.

An NSService can also be a group of other NSServices. An NSService that is a group can be of the following types:

- Layer 2
- Layer 3 and above

You cannot change the type after you create an NSService. Some NSServices are predefined. You cannot modify or delete them.

Create an NSService in Manager Mode

You can create an NSService to specify the characteristics that network matching uses, or to define the type of traffic to block or allow in firewall rules.

Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Select **Inventory > Services > Add**.
3. Enter a name.
4. (Optional) Enter a description.
5. Select **Specify a protocol** to configure an individual service, or select **Group existing services** to configure a group of NSServices.
6. For an individual service, select a type of service and a protocol.

   The available types are **Ether**, **IP**, **IGMP**, **ICMP**, **ALG**, and **L4 Port Set**
7. For a service group, select a type and members for the group.

   The available types are **Layer 2** and **Layer 3 and above**.
8. Click **ADD**.

Manage Tags for a VM in Manager Mode

You can see the list of VMs in the inventory. You can also add tags to a VM to make searching easier.

Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Inventory > Virtual Machines** from the navigation panel.

The list of VMs is displayed with 4 columns: Virtual Machine, External ID, Source, and Tag. Click the filter icon in the first three columns' heading to filter the list. Enter a string of characters to do a partial match. If the string in the column contains the string that you entered, the entry is displayed. Enter a string of characters enclosed in double quotes to do an exact match. If the string in the column exactly matches the string that you entered, the entry is displayed.

3 Select **Inventory > Virtual machines** from the navigation panel.

4 Select a VM.

5 Click **MANAGE TAGS**.

6 Add or delete tags.

Option	Action
Add a tag	Click **ADD** to specify a tag and optionally a scope.
Delete a tag	Select an existing tag and click **DELETE**.

The maximum number of tags that can be assigned from the NSX Manager to a virtual machine is 25. The maximum number of tags for all other managed objects such as logical switches or ports, is 30.

Note Tags used in system groups must not be assigned to regular VMs. For example, if tag **ABC** is used in the membership criteria of the Edge_NSGroup system group, this tag must not be assigned to any regular VM.

7 Click **Save**.

DHCP in Manager Mode

You can configure DHCPv4 in **Manager** mode.

You cannot configure or modify a DHCPv6 server configuration in **Manager** mode. You must use any of the following to configure or modify a DHCPv6 server:

- Policy mode
- Policy API

- Manager API

Note If you use **Manager** mode to modify objects created in the **Policy** mode, some settings might not be configurable. These read-only settings have this icon next to them: ⊖. See Chapter 1 NSX Manager for more information.

DHCP

DHCP (Dynamic Host Configuration Protocol) allows clients to automatically obtain network configuration, such as IP address, subnet mask, default gateway, and DNS configuration, from a DHCP server.

You can create DHCP servers to handle DHCP requests and create DHCP relay services to relay DHCP traffic to external DHCP servers. However, you should not configure a DHCP server on a logical switch and also configure a DHCP relay service on a router port that the same logical switch is connected to. In such a scenario, DHCP requests will only go to the DHCP relay service.

If you configure DHCP servers, to improve security, configure a DFW rule to allow traffic on UDP ports 67 and 68 only for valid DHCP server IP addresses.

To block DHCP packets for ports 67 and 68 configure a DFW rule with the following:

Source	Destination	Service	Rule
ANY	ANY	ANY	BLOCK

To allow DHCP packets configure a DFW rule with the following:

Source	Destination	Service	Rule
ANY	ANY	ports 67 and 68, TCP	ALLOW

Note In this release, the DHCP server does not support guest VLAN tagging.

Create a DHCP Server Profile in Manager Mode

A DHCP server profile specifies an NSX Edge cluster or members of an NSX Edge cluster. A DHCP server with this profile services DHCP requests from VMs on logical switches that are connected to the NSX Edge nodes that are specified in the profile.

Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Networking > DHCP > Server Pofiles > Add**.

3 Enter a name and optional description.

4 Select an NSX Edge cluster from the drop-down menu.

5 (Optional) Select members of the NSX Edge cluster.

You can specify up to 2 members.

What to do next

Create a DHCP server. See Create a DHCP Server in Manager Mode.

Create a DHCP Server in Manager Mode

You can create DHCP servers to service DHCP requests from VMs that are connected to logical switches.

Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Networking > DHCP > Servers > Add**.

3 Enter a name and optional description.

4 Enter the IP address of the DHCP server and its subnet mask in CIDR format.

For example, enter `192.168.1.2/24`.

5 (Required) Select a DHCP profile from the drop-down menu.

6 (Optional) Enter common options such as domain name, default gateway, DNS servers, and subnet mask.

7 (Optional) Enter classless static route options.

8 (Optional) Enter other options.

9 Click **Save**.

10 Select the newly created DHCP server.

11 Expand the IP Pools section.

12 Click **Add** to add IP ranges, default gateway, lease duration, warning threshold, error threshold, classless static route option, and other options.

13 Expand the Static Bindings section.

14 Click **Add** to add static bindings between MAC addresses and IP addresses, default gateway, hostname, lease duration, classless static route option, and other options.

What to do next

Attach a DHCP server to a logical switch. See Attach a DHCP Server to a Logical Switch in Manager Mode.

Attach a DHCP Server to a Logical Switch in Manager Mode

You must attach a DHCP server to a logical switch before the DHCP server can process DHCP requests from VMs connected to the switch.

Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Select **Networking > Logical Switches > Switches**.
   a. Select a logical switch.
   b. Click **Actions > Attach to a DHCP Server**.
3. Alternatively, select **Networking > DHCP > Servers** .
   a. Select a DHCP server.
   b. Click **Actions > Attach to Logical Switch**.

Detach a DHCP Server from a Logical Switch in Manager Mode

You can detach a DHCP server from a logical switch to reconfigure your environment.

Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Select **Networking > Logical Switches**.
3. Click the logical switch that you intend to detach a DHCP server from.
4. Click **Actions > Detach from the DHCP Server**.

Create a DHCP Relay Profile in Manager Mode

A DHCP relay profile specifies one or more external DHCP or DHCPv6 servers. When you create a DHCP/DHCPv6 relay service, you must specify a DHCP relay profile.

Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Networking > DHCP > Relay Profiles > Add**.

3 Enter a name and optional description.

4 Enter one or more external DHCP/DHCPv6 server addresses.

What to do next

Create a DHCP/DHCPv6 relay service. See Create a DHCP Relay Service in Manager Mode.

Create a DHCP Relay Service in Manager Mode

You can create a DHCP relay service to relay traffic between DHCP clients and DHCP servers that are not created in NSX.

Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Networking > DHCP > Relay Services > Add**.

3 Enter a name and optional description.

4 Select a DHCP relay profile from the drop-down menu.

What to do next

Add a DHCP service to a logical router port. See Add a DHCP Relay Service to a Logical Router Port in Manager Mode.

Add a DHCP Relay Service to a Logical Router Port in Manager Mode

You can add a DHCP relay service to a logical router port. VMs on the logical switch that is attached to that port can communicate with the DHCP servers that are configured in the relay service.

Prerequisites

- Verify you have a configured DHCP relay service. See Create a DHCP Relay Service in Manager Mode.
- Verify that the router port is of type **Downlink**.
- Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Select **Networking > Tier-0 Logical Routers**.
3. Select the appropriate router to display more information and configuration options.
4. Select **Configuration > Router Ports**.
5. Select the router port that connects to the desired logical switch and click **Edit**.
6. Select a DHCP relay service from the **Relay Service** drop-down list and click **Save**.

   You can also select a DHCP relay service when you add a new logical router port.

Delete a DHCP Lease

In some situations, you might want to delete a DHCP lease. For example, if you want a DHCP client to get a different IP address, or if a client shuts down without releasing its IP address and you want the address to be available to other clients.

You can use the following API to delete a DHCP lease:

```
DELETE /api/v1/dhcp/servers/<server-id>/leases?ip=<ip>&mac=<mac>
```

To ensure that the correct lease is deleted, call the following API before and after the `DELETE` API:

```
GET /api/v1/dhcp/servers/<server-id>/leases
```

After calling the `DELETE` API, make sure that the output of the `GET` API does not show the lease that was deleted.

For more information, see the *NSX API Guide*.

Metadata Proxies

With a metadata proxy server, VM instances can retrieve instance-specific metadata from an OpenStack Nova API server.

The following steps describe how a metadata proxy works:

1. A VM sends an HTTP GET to http://169.254.169.254:80 to request some metadata.

2 The metadata proxy server that is connected to the same logical switch as the VM reads the request, makes appropriate changes to the headers, and forwards the request to the Nova API server.

3 The Nova API server requests and receives information about the VM from the Neutron server.

4 The Nova API server finds the metadata and sends it to the metadata proxy server.

5 The metadata proxy server forwards the metadata to the VM.

A metadata proxy server runs on an NSX Edge node. For high availability, you can configure metadata proxy to run on two or more NSX Edge nodes in an NSX Edge cluster.

Add a Metadata Proxy Server in Manager Mode

A metadata proxy server enables VMs to retrieve metadata from an OpenStack Nova API server.

Prerequisites

- Verify that you have created an NSX Edge cluster. For more information, see *NSX Installation Guide*.
- Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Networking > DHCP > Metadata Proxies > Add** .

3 Enter a name for the metadata proxy server.

4 (Optional) Enter a description.

5 Enter the URL and port for the Nova server.

The valid port range is 3000 - 9000.

6 Enter a value for **Secret**.

7 Select an NSX Edge cluster from the drop-down list.

8 (Optional) Select members of the NSX Edge cluster.

What to do next

Attach the metadata proxy server to a logical switch.

Attach a Metadata Proxy Server to a Logical Switch in Manager Mode

To provide metadata proxy services to VMs that are connected to a logical switch, you must attach a metadata proxy server to the switch.

Prerequisites

- Verify that you have created a logical switch. For more information, see Create a Logical Switch in Manager Mode.
- Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Networking > DHCP > Metadata Proxies**.

3 Select a metadata proxy server.

4 Select the menu option **Actions > Attach to Logical Switch**

5 Select a logical switch from the drop-down list.

Results

You can also attach a metadata proxy server to a logical switch by navigating to **Networking > Logical Switches > Switches**, selecting a switch, and selecting the menu option **Actions > Add to a Metadata Proxy**.

Detach a Metadata Proxy Server from a Logical Switch in Manager Mode

To stop providing metadata proxy services to VMs that are connected to a logical switch or use a different metadata proxy server, you can detach a metadata proxy server from a logical switch.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Networking > DHCP > Metadata Proxies**.

3 Select a metadata proxy server.

4 Select the menu option **Actions > Detach from Logical Switch**

5 Select a logical switch from the drop-down list.

Results

You can also detach a metadata proxy server from a logical switch by navigating to **Networking > Logical Switches > Switches**, selecting a switch, and selecting the menu option **Actions > Detach from the Metadata Proxy**.

IP Address Management in Manager Mode

With IP address management (IPAM), you can create IP blocks to support NSX Container Plugin (NCP). For more info about NCP, see the *NSX Container Plug-in for Kubernetes and Tanzu Application Service - Installation and Administration Guide*.

Note If you use **Manager** mode to modify objects created in the **Policy** mode, some settings might not be configurable. These read-only settings have this icon next to them: ⊖. See Chapter 1 NSX Manager for more information.

Manage IP Blocks in Manager Mode

Setting up NSX Container Plugin requires that you create IP blocks for the containers.

Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Select **Networking > IP Address Pools > IP Blocks**.
3. To add an IP block, click **Add**.
   a. Enter a name and optionally a description.
   b. Enter an IP block in CIDR format. For example, 10.10.10.0/24.
4. To edit an IP block, click the name of an IP block.
   a. In the **Overview** tab, click **Edit**.
      You can change the name, description, or the IP block value.
5. To manage the tags of an IP block, click the name of an IP block.
   a. In the **Overview** tab, click **Manage**.
      You can add or delete tags.
6. To delete one or more IP blocks, select the blocks.
   a. Click **Delete**.
      You cannot delete an IP block that has its subnet allocated.

Manage Subnets for IP Blocks in Manager Mode

You can add or delete subnets for IP blocks.

Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Networking > IP Address Pools > IP Blocks**.

3 Click the name of an IP block.

4 Click the **Subnets** tab.

5 To add a subnet, click **Add**.

   a Enter a name and optionally a description.

   b Enter the size of the subnet.

6 To delete one or more subnets, select the subnets.

   a Click **Delete**.

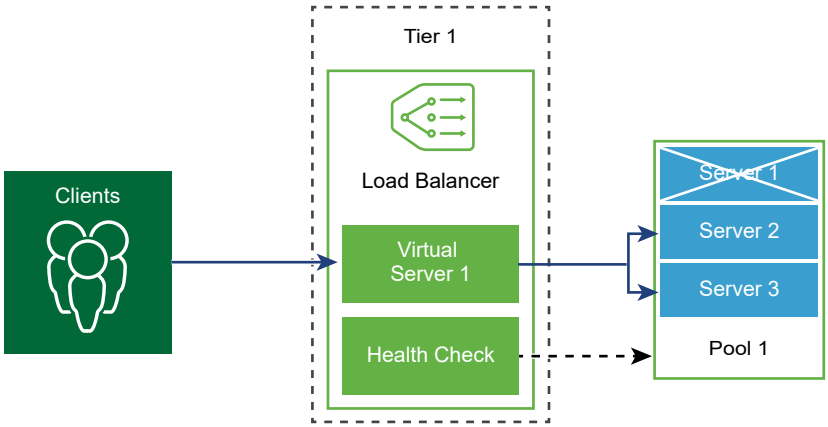
Load Balancing in Manager Mode

This information covers the NSX load balancing configuration in **Manager** mode.

.

Note If you use **Manager** mode to modify objects created in the **Policy** mode, some settings might not be configurable. These read-only settings have this icon next to them: ⊖. See Chapter 1 NSX Manager for more information.

The NSX logical load balancer offers high-availability service for applications and distributes the network traffic load among multiple servers.



The load balancer distributes incoming service requests evenly among multiple servers in such a way that the load distribution is transparent to users. Load balancing helps in achieving optimal resource utilization, maximizing throughput, minimizing response time, and avoiding overload.

You can map a virtual IP address to a set of pool servers for load balancing. The load balancer accepts TCP, UDP, HTTP, or HTTPS requests on the virtual IP address and decides which pool server to use.

Depending on your environment needs, you can scale the load balancer performance by increasing the existing virtual servers and pool members to handle heavy network traffic load.

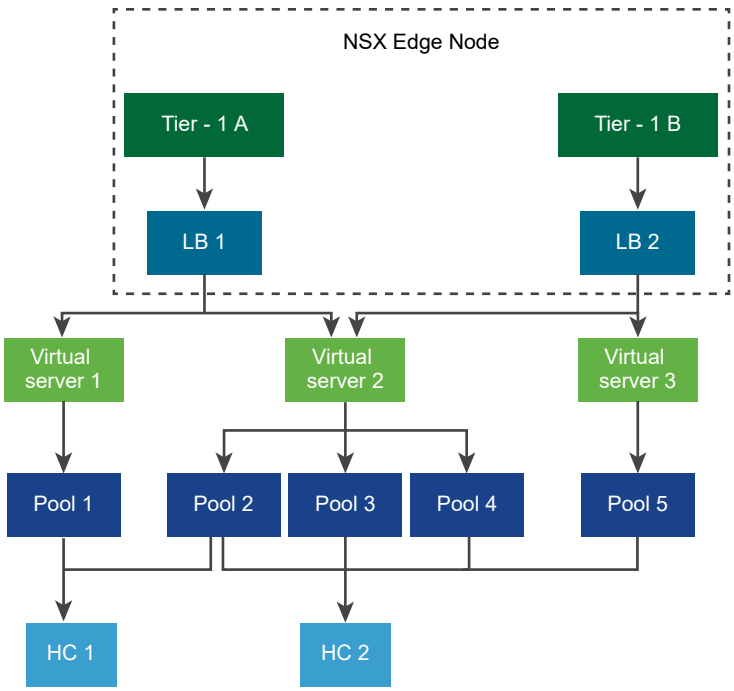
Note Logical load balancer is supported only on the Tier-1 logical router. One load balancer can be attached only to a Tier-1 logical router.

Key Load Balancer Concepts

Load balancer includes virtual servers, server pools, and health checks monitors.

A load balancer is connected to a Tier-1 logical router. The load balancer hosts single or multiple virtual servers. A virtual server is an abstract of an application service, represented by a unique combination of IP, port, and protocol. The virtual server is associated to single to multiple server pools. A server pool consists of a group of servers. The server pools include individual server pool members.

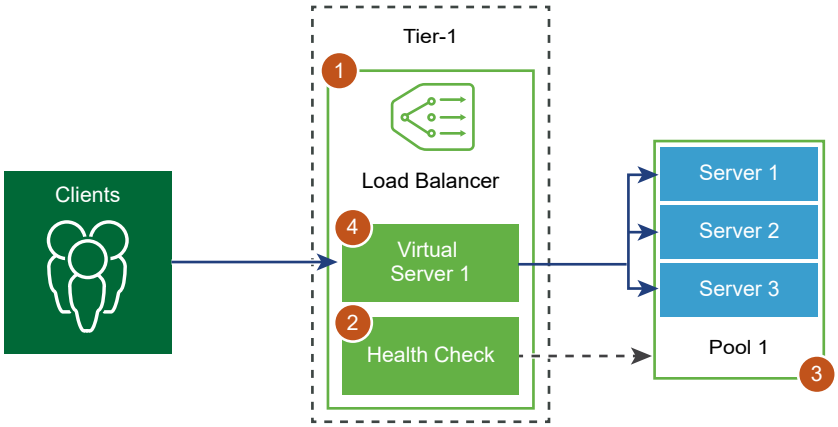
To test whether each server is correctly running the application, you can add health check monitors that check the health status of a server.



Configuring Load Balancer Components

To use logical load balancers, you must start by configuring a load balancer and attaching it to a Tier-1 logical router.

Next, you can set up health check monitoring for your servers. You must then configure server pools for your load balancer. Finally, you must create a layer 4 or layer 7 virtual server for your load balancer.

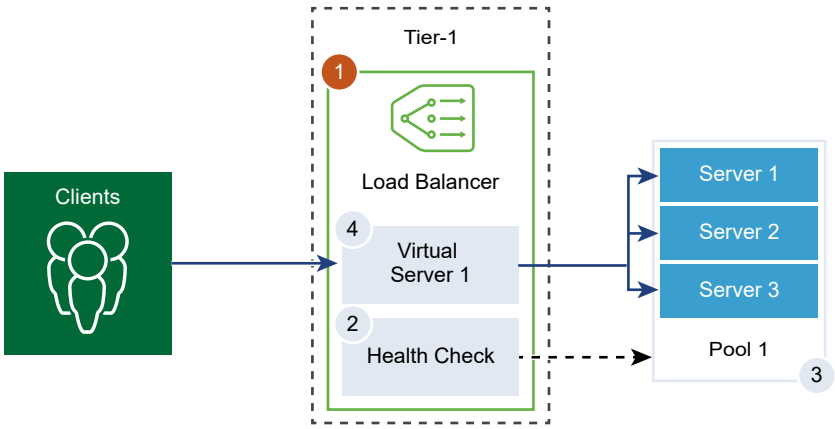


Create a Load Balancer in Manager Mode

Load balancer is created and attached to the Tier-1 logical router.

Important The information in this topic is specific to administering your environment in manager mode. For more information about manager mode and policy mode, see Chapter 1 NSX Manager. For information about load balancers in policy mode, see Chapter 10 Load Balancer .

You can configure the level of error messages you want the load balancer to add to the error log. Avoid setting the log level to DEBUG on load balancers with significant traffic due to the number of messages printed to the log that affect performance.



Prerequisites

- Verify that a Tier-1 logical router is configured. See Create a Tier-1 Logical Router in Manager Mode .
- Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Networking > Load Balancer > Add**.

3 Enter a name and a description for the load balancer.

4 Select the load balancer virtual server size and number of pool members based on your available resources.

5 Define the severity level of the error log from the drop-down menu.

Load balancer collects information about encountered issues of different severity levels to the error log.

6 Click **OK**.

7 Associate the newly created load balancer to a virtual server.

a Select the load balancer and click **Actions > Attach to a Virtual Server**.

b Select an existing virtual server from the drop-down menu.

c Click **OK**.

8 Attach the newly created load balancer to a Tier-1 logical router.

a Select the load balancer and click **Actions > Attach to a Logical Router**.

b Select an existing Tier-1 logical router from the drop-down menu.

The Tier-1router must be in the Active-Standby mode.

c Click **OK**.

9 (Optional) Delete the load balancer.

If you no longer want to use this load balancer, you must first detach the load balancer from the virtual server and Tier-1 logical router.

Configure an Active Health Monitor in Manager Mode

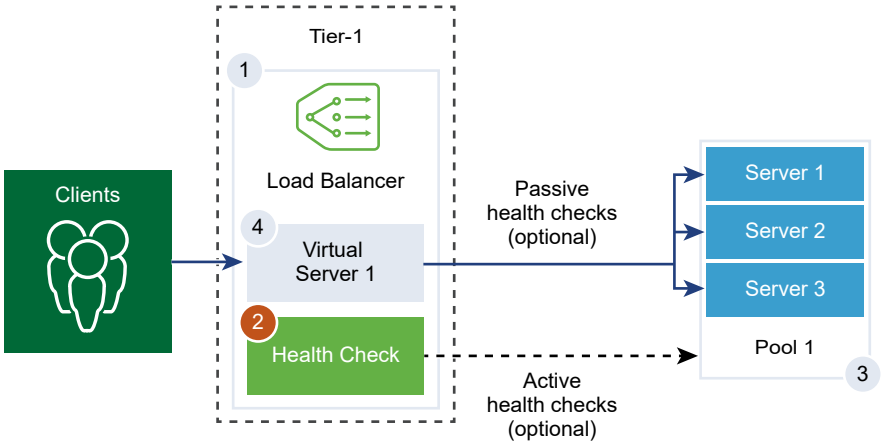
The active health monitor is used to test whether a server is available. The active health monitor uses several types of tests such as sending a basic ping to servers or advanced HTTP requests to monitor the application health.

Servers that fail to respond within a certain time period or respond with errors are excluded from future connection handling until a subsequent periodic health check finds these servers to be healthy.

Active health checks are performed on server pool members after the pool member is attached to a virtual server and that virtual server is attached to a Tier-1 gateway (previously called a Tier-1 logical router).

If the Tier-1 gateway is connected to a Tier-0 gateway, a router link port is created and its IP address (typically in the `100.64.x.x` format) is used to perform the health check for the load balancer service. If the Tier-1 gateway is standalone (has only one centralized service port and is not connected to a Tier-0 gateway), the centralized service port IP address is used to perform the health check for the load balancer service. See Create a Standalone Tier-1 Logical Router in Manager Mode for information about standalone Tier-1 gateways.

Note More than one active health monitor can be configured per server pool.



Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Networking > Load Balancing > Monitors > Active Health Monitors > Add**.

3 Enter a name and description for the active health monitor.

4 Select a health check protocol for the server from the drop-down menu.

You can also use predefined protocols in NSX Manager; `http-monitor`, `https-monitor`, `Icmp-monitor`, `Tcp-monitor`, and `Udp-monitor`.

5 Set the value of the monitoring port.

6 Configure the values to monitor a service pool.

You can also accept the default active health monitor values.

Option	Description
Monitoring Interval	Set the time in seconds that the monitor sends another connection request to the server.
Fall Count	Set a value when the consecutive failures reach this value, the server is considered temporarily unavailable.
Rise Count	Set a number after this timeout period, the server is tried again for a new connection to see if it is available.
Timeout Period	Set the number of times the server is tested before it is considered as DOWN.

For example, if the monitoring interval is set as 5 seconds and the timeout as 15 seconds, the load balancer send requests to the server every 5 seconds. In each probe, if the expected response is received from the server within 15 seconds, then the health check result is OK. If not, then the result is CRITICAL. If the recent three health check results are all UP, the server is considered as UP.

7 If you select HTTP as the health check protocol, complete the following details.

Option	Description
HTTP Method	Select the method for detecting the server status from the drop-down menu, GET, OPTIONS, POST, HEAD, and PUT.
HTTP Request URL	Enter the request URI for the method. ASCII control characters (backspace, vertical tab, horizontal tab, line feed, etc), unsafe characters such as a `space`, `\`, `<`, `>`, `{`, `}`, and any character outside the ASCII character set are not allowed in the request URL and should be encoded. For example, replace a space with a plus (+) sign, or with %20.
HTTP Request Version	Select the supported request version from the drop-down menu. You can also accept the default version, HTTP_VERSION_1_1.
HTTP Request Body	Enter the request body. Valid for the POST and PUT methods.
HTTP Response Code	Enter the string that the monitor expects to match in the status line of HTTP response body. The response code is a comma-separated list. For example, 200,301,302,401.
HTTP Response Body	If the HTTP response body string and the HTTP health check response body match, then the server is considered as healthy.

8 If you select HTTPs as the health check protocol, complete the following details.

a Select the SSL protocol list.

TLS versions TLS1.1 and TLS1.2 versions are supported and enabled by default. TLS1.0 is supported, but disabled by default.

b Click the arrow and move the protocols into the selected section.

c Assign a default SSL cipher or create a custom SSL cipher.

d Complete the following details for HTTP as the health check protocol.

Option	Description
HTTP Method	Select the method for detecting the server status from the drop-down menu: `GET`, `OPTIONS`, `POST`, `HEAD`, and `PUT`.
HTTP Request URL	Enter the request URI for the method. ASCII control characters (backspace, vertical tab, horizontal tab, line feed, etc), unsafe characters such as a `space`, `\`, `<`, `>`, `{`, `}`, and any character outside the ASCII character set are not allowed in the request URL and should be encoded. For example, replace a space with a plus (+) sign, or with %20.
HTTP Request Version	Select the supported request version from the drop-down menu. You can also accept the default version, HTTP_VERSION_1_1.
HTTP Request Body	Enter the request body. Valid for the POST and PUT methods.
HTTP Response Code	Enter the string that the monitor expects to match in the status line of HTTP response body. The response code is a comma-separated list. For example, 200,301,302,401.
HTTP Response Body	If the HTTP response body string and the HTTP health check response body match, then the server is considered as healthy.

9 If you select ICMP as the health check protocol, assign the data size in byte of the ICMP health check packet.

10 If you select TCP as the health check protocol, you can leave the parameters empty.

If both the sent and expected are not listed, then a three-way handshake TCP connection is established to validate the server health. No data is sent. Expected data if listed has to be a string and can be anywhere in the response. Regular expressions are not supported.

11 If you select UDP as the health check protocol, complete the following required details.

Required Option	Description
UDP Data Sent	Enter the string to be sent to a server after a connection is established.
UDP Data Expected	Enter the string expected to receive from the server. Only when the received string matches this definition, is the server is considered as UP.

12 Click **Finish**.

What to do next

Associate the active health monitor with a server pool. See Add a Server Pool for Load Balancing in Manager Mode.

Configure Passive Health Monitors in Manager Mode

Load balancers perform passive health checks to monitor failures during client connections and mark servers causing consistent failures as DOWN.

Passive health check monitors client traffic going through the load balancer for failures. For example, if a pool member sends a TCP Reset (RST) in response to a client connection, the load balancer detects that failure. If there are multiple consecutive failures, then the load balancer considers that server pool member to be temporarily unavailable and stops sending connection requests to that pool member for some time. After some time, the load balancer sends a connection request to check if the pool member has recovered. If that connection is successful, then the pool member is considered healthy. Otherwise, the load balancer waits for some time and tries again.

Passive health check considers the following scenarios to be failures in client traffic.

- For server pools associated with Layer 7 virtual servers, if the connection to the pool member fails. For example, if the pool member sends a TCP RST when the load balancer tries to connect or perform a SSL handshake between load balancer and the pool member fails.
- For server pools associated with Layer 4 TCP virtual servers, if the pool member sends a TCP RST in response to client TCP SYN or does not respond at all.
- For server pools associated with Layer 4 UDP virtual servers, if a port is unreachable or a destination unreachable ICMP error message is received in response to a client UDP packet.

Server pools associated to Layer 7 virtual servers, the failed connection count is incremented when any TCP connection errors, for example, TCP RST failure to send data or SSL handshake failures occur.

Server pools associated to Layer 4 virtual servers, if no response is received to a TCP SYN sent to the server pool member or if a TCP RST is received in response to a TCP SYN, then the server pool member is considered as DOWN. The failed count is incremented.

For Layer 4 UDP virtual servers, if an ICMP error such as, port or destination unreachable message is received in response to client traffic, then it is considered as DOWN.

Note One passive health monitor can be configured per server pool.

Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Networking > Load Balancing > Monitors > Passive Health Monitors > Add**.

3 Enter a name and description for the passive health monitor.

4 Configure the values to monitor a service pool.

You can also accept the default active health monitor values.

Option	Description
Fall Count	Set a value when the consecutive failures reach this value, the server is considered temporarily unavailable.
Timeout Period	Set the number of times the server is tested before it is considered as DOWN.

For example, when the consecutive failures reach the configured value 5, that member is considered temporarily unavailable for 5 seconds. After this period, that member is tried again for a new connection to see if it is available. If that connection is successful, then the member is considered available and the failed count is set to zero. However, if that connection fails, then it is not used for another timeout interval of 5 seconds.

5 Click **OK**.

What to do next

Associate the passive health monitor with a server pool. See Add a Server Pool for Load Balancing in Manager Mode.

Add a Server Pool for Load Balancing in Manager Mode

Server pool consists of one or more servers that are configured and running the same application. A single pool can be associated to both Layer 4 and Layer 7 virtual servers.

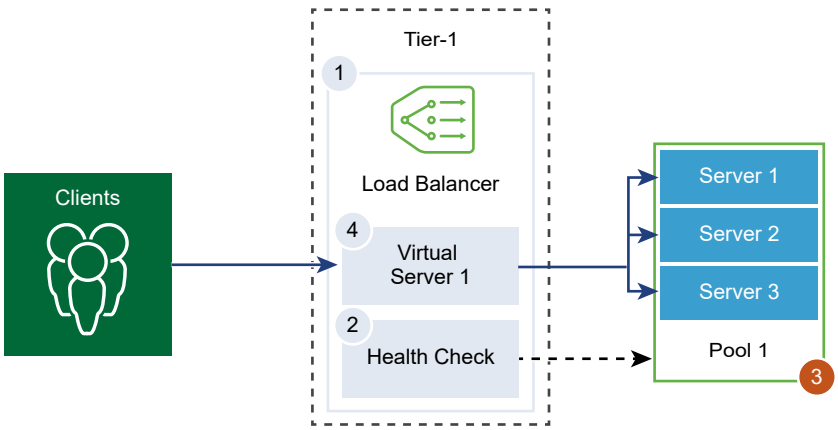
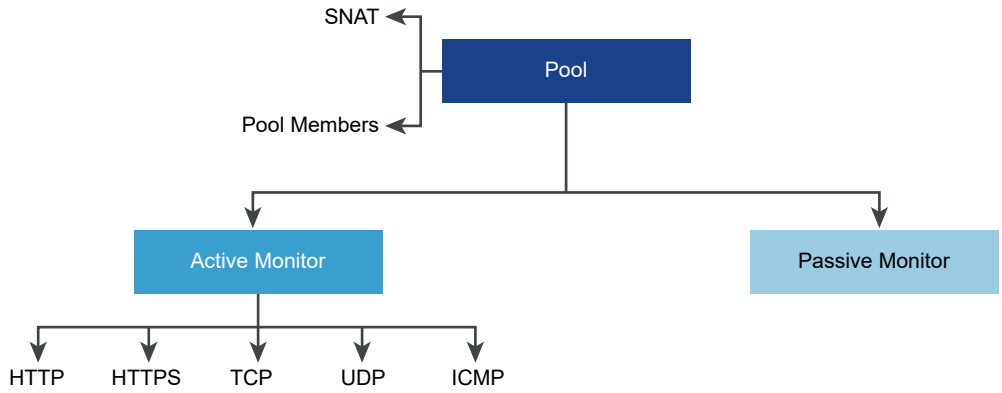


Figure 25-13. Server Pool Parameter Configuration



Prerequisites

- If you use dynamic pool members, an NSGroup must be configured. See Create an NSGroup in Manager Mode.
- Depending on the monitoring you use, verify that active or passive health monitors are configured. See Configure an Active Health Monitor in Manager Mode or Configure Passive Health Monitors in Manager Mode.
- Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Select **Networking > Load Balancing > Server Pools > Add**.
3. Enter a name and description for the load balancer pool.

   You can optionally describe the connections managed by the server pool.
4. Select the algorithm balancing method for the server pool.

   Load balancing algorithm controls how the incoming connections are distributed among the members. The algorithm can be used on a server pool or a server directly.

   All load balancing algorithms skip servers that meet any of the following conditions:

   - Admin state is set to DISABLED.
   - Admin state is set to GRACEFUL_DISABLED and no matching persistence entry.
   - Active or passive health check state is DOWN.

- Connection limit for the maximum server pool concurrent connections is reached.

Option	Description
ROUND_ROBIN	Incoming client requests are cycled through a list of available servers capable of handling the request. Ignores the server pool member weights even if they are configured.
WEIGHTED_ROUND_ROBIN	Each server is assigned a weight value that signifies how that server performs relative to other servers in the pool. The value determines how many client requests are sent to a server compared to other servers in the pool. This load balancing algorithm focuses on fairly distributing the load among the available server resources.
LEAST_CONNECTION	Distributes client requests to multiple servers based on the number of connections already on the server. New connections are sent to the server with the fewest connections. Ignores the server pool member weights even if they are configured.
WEIGHTED_LEAST_CONNECTION	Each server is assigned a weight value that signifies how that server performs relative to other servers in the pool. The value determines how many client requests are sent to a server compared to other servers in the pool. This load balancing algorithm focuses on using the weight value to distribute the load among the available server resources fairly. By default, the weight value is 1 if the value is not configured and slow start is enabled.
IP-HASH	Selects a server based on a hash of the source IP address and the total weight of all the running servers.

5 Toggle the TCP Multiplexing button to enable this menu item.

With TCP multiplexing, you can use the same TCP connection between a load balancer and the server for sending multiple client requests from different client TCP connections.

6 Set the maximum number of TCP multiplexing connections per pool that are kept alive to send future client requests.

7 Select the Source NAT (SNAT) mode.

Depending on the topology, SNAT might be required so that the load balancer receives the traffic from the server destined to the client. SNAT can be enabled per server pool.

Translation Mode	Description
Transparent	Load balancer uses the client IP address and port spoofing while establishing connections to the servers. SNAT is not required.
Auto Map	Load Balancer uses the interface IP address and ephemeral port to continue the communication with a client initially connected to one of the server's established listening ports. SNAT is required. Enable port overloading to allow the same SNAT IP and port to be used for multiple connections if the tuple (source IP, source port, destination IP, destination port, and IP protocol) is unique after the SNAT process is performed. You can also set the port overload factor to allow the maximum number of times a port can be used simultaneously for multiple connections.
IP List	Specify a single IP address range, for example, 1.1.1.1-1.1.1.10 to be used for SNAT while connecting to any of the servers in the pool. By default, from 4000 through 64000 port range is used for all configured SNAT IP addresses. Port ranges from 1000 through 4000 are reserved for purposes such as, health checks and connections initiated from Linux applications. If multiple IP addresses are present, then they are selected in a Round Robin manner. Enable port overloading to allow the same SNAT IP and port to be used for multiple connections if the tuple (source IP, source port, destination IP, destination port, and IP protocol) is unique after the SNAT process is performed. You can also set the port overload factor to allow the maximum number of times a port can be used simultaneously for multiple connections.

8 Select the server pool members.

Server pool consists of single or multiple pool members. Each pool member has an IP address and a port.

Each server pool member can be configured with a weight for use in the load balancing algorithm. The weight indicates how much more or less load a given pool member can handle relative to other members in the same pool.

Designating a pool member as a backup member works with the health monitor to provide an active/standby state. If active members fail a health check, traffic failover occurs for backup members.

Membership Type	Description
Static	Click **Add** to include a static pool member. You can also clone an existing static pool member.
Dynamic	Select the NSGroup from the drop-down menu. The server pool membership criteria is defined in the group. You can optionally, define the maximum group IP address list.

9 Enter the minimum number of active members the server pool must always maintain.

10 Select an active and passive health monitor for the server pool from the drop-down menu.

Setting an active and passive health monitor for the server pool is optional. When you select an active health monitor and if the Tier-1 gateway is connected to a Tier-0 gateway, a router link port is created. The router link port's IP address (typically in the `100.64.x.x` format) is used to perform the health check for the load balancer service. If the Tier-1 gateway is standalone (has only one centralized service port and is not connected to a Tier-0 gateway), the centralized service port IP address is used to perform the health check for the load balancer service. See Create a Standalone Tier-1 Logical Router in Manager Mode for information about standalone Tier-1 gateways.

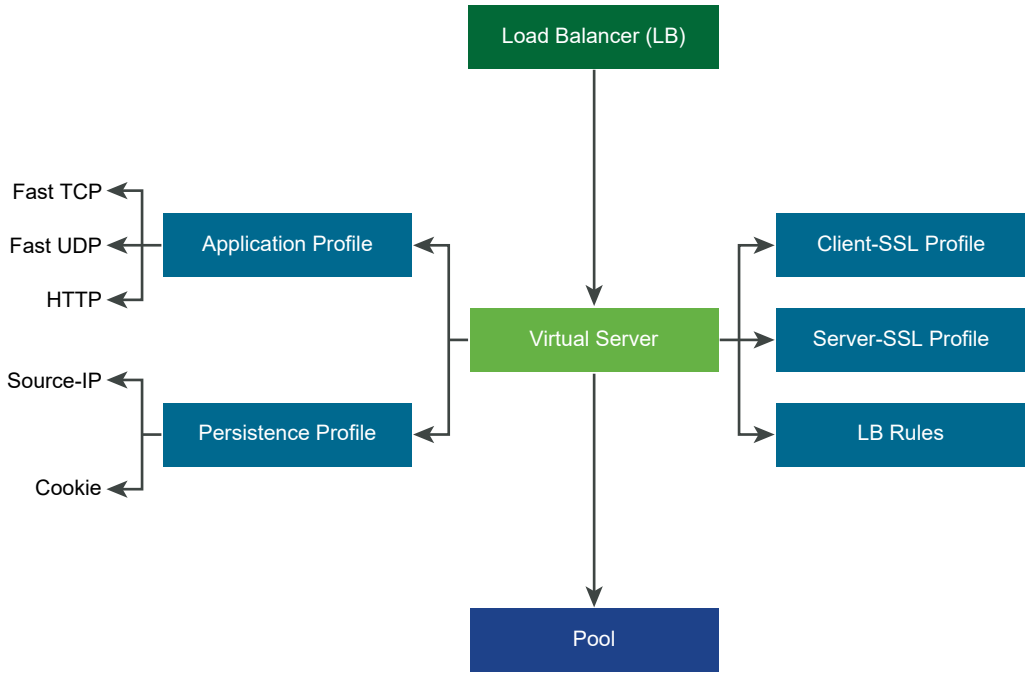
Add a firewall rule to allow the IP address to perform the health check for the load balancer service.

11 Click **Finish.**

Configuring Virtual Server Components

With the virtual server there are several components that you can configure such as application profiles, persistent profiles, and load balancer rules.

Figure 25-14. Virtual Server Components

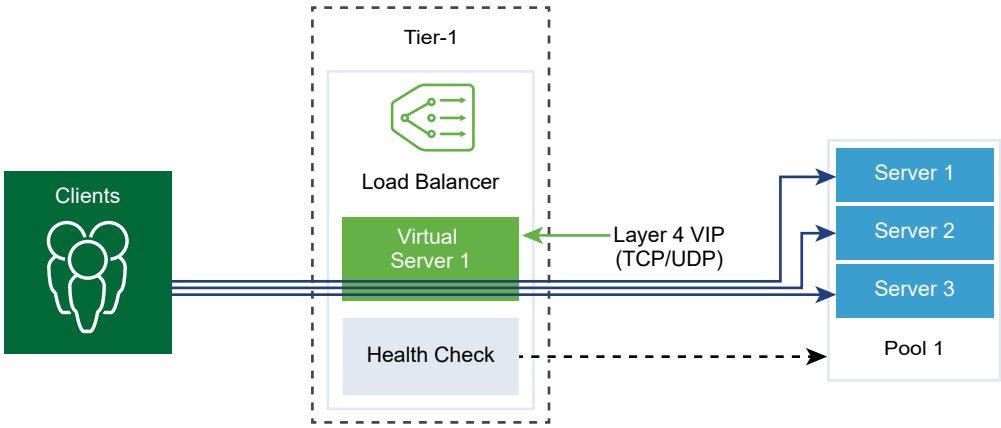


Configure Application Profiles in Manager Mode

Application profiles are associated with virtual servers to enhance load balancing network traffic and simplify traffic-management tasks.

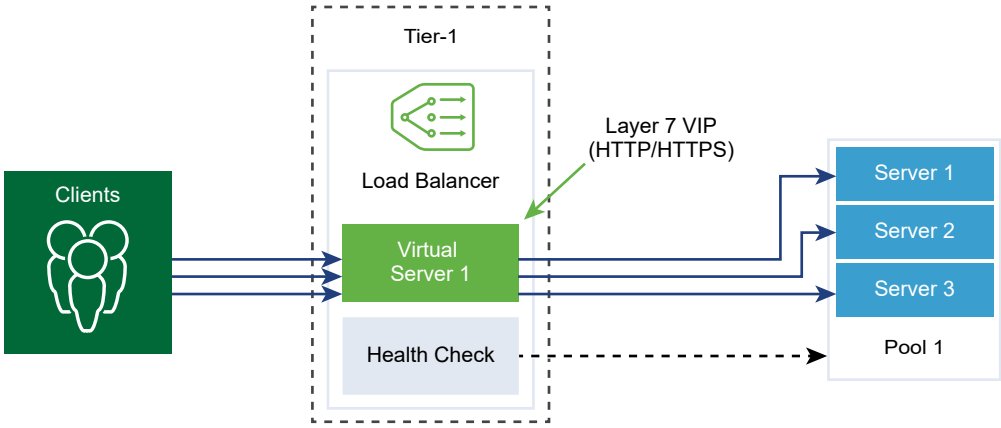
Application profiles define the behavior of a particular type of network traffic. The associated virtual server processes network traffic according to the values specified in the application profile. Fast TCP, Fast UDP, and HTTP application profiles are the supported types of profiles.

Figure 25-15. Layer 4 TCP and UDP Application Profile



TCP application profile is used by default when no application profile is associated to a virtual server. TCP and UDP application profiles are used when an application is running on a TCP or UDP protocol and does not require any application level load balancing such as, HTTP URL load balancing. These profiles are also used when you only want Layer 4 load balancing, which has faster performance and supports connection mirroring.

Figure 25-16. Layer 7 HTTPS Application Profile



HTTP application profile is used for both HTTP and HTTPS applications when the load balancer needs to take actions based on Layer 7 such as, load balancing all images requests to a specific server pool member or terminating HTTPS to offload SSL from pool members. Unlike the TCP application profile, the HTTP application profile terminates the client TCP connection before selecting the server pool member.

Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Networking > Load Balancing > Profiles > Application Profiles**.

3 Create a Fast TCP application profile.

   a Select **Add > Fast TCP Profile** from the drop-down menu.

   b Enter a name and a description for the Fast TCP application profile.

c Complete the application profile details.

You can also accept the default FAST TCP profile settings.

Option	Description
Connection Idle Timeout	Enter the time in seconds on how long the server can remain idle after a TCP connection is established. Set the idle time to the actual application idle time and add a few more seconds so that the load balancer does not close its connections before the application does.
Connection Close Timeout	Enter the time in seconds that the TCP connection both FINs or RST must be kept for an application before closing the connection. A short closing timeout might be required to support fast connection rates.
HA Flow Mirroring	Toggle the button to make all the flows to the associated virtual server mirrored to the HA standby node.

d Click **OK**.

4 Create a Fast UDP application profile.

You can also accept the default UDP profile settings.

a Select **Add > Fast UDP Profile** from the drop-down menu.

b Enter a name and a description for the Fast UDP application profile.

c Complete the application profile details.

Option	Description
Idle Timeout	Enter the time in seconds on how long the server can remain idle after a UDP connection is established. UDP is a connectionless protocol. For load balancing purposes, all the UDP packets with the same flow signature such as, source and destination IP address or ports and IP protocol received within the idle timeout period are considered to belong to the same connection and sent to the same server. If no packets are received during the idle timeout period, the connection which is an association between the flow signature and the selected server is closed.
HA Flow Mirroring	Toggle the button to make all the flows to the associated virtual server mirrored to the HA standby node.

d Click **OK**.

5 Create an HTTP application profile.

You can also accept the default HTTP profile settings.

HTTP application profile is used for both HTTP and HTTPS applications.

a Select **Add > Fast HTTP Profile** from the drop-down menu.

b Enter a name and a description for the HTTP application profile.

c Complete the application profile details.

Option	Description
Name and Description	Enter a name and a description for the HTTP application profile.
Idle Timeout	Enter the time in seconds on how long client idle connections remain before the load balancer closes them (FIN).
Request Header Size	Specify the maximum buffer size in bytes used to store HTTP request headers.
Response Header Size	Specify the maximum buffer size in bytes used to store HTTP response headers. The default is 4096, and the maximum is 65536.
Redirection	■ None - If a website is temporarily down, user receives a page not found error message. ■ HTTP Redirect - If a website is temporarily down or has moved, incoming requests for that virtual server can be temporarily redirected to a URL specified here. Only a static redirection is supported. For example, if HTTP Redirect is set to http://sitedown.abc.com/sorry.html, then irrespective of the actual request, for example, http://original_app.site.com/home.html or http://original_app.site.com/somepage.html, incoming requests are redirected to the specified URL when the original website is down. ■ HTTP to HTTPS Redirect - Certain secure applications might want to force communication over SSL, but instead of rejecting non-SSL connections, they can redirect the client request to use SSL. With HTTP to HTTPS Redirect, you can preserve both the host and URI paths and redirect the client request to use SSL. For HTTP to HTTPS redirect, the HTTPS virtual server must have port 443 and the same virtual server IP address must be configured on the same load balancer. For example, a client request for http://app.com/path/page.html is redirected to https://app.com/path/page.html. If either the host name or the URI must be modified while redirecting, for example, redirect to https://secure.app.com/path/page.html, then load balancing rules must be used.
Tags	Enter tags to make searching easier. You can specify a tag to set a scope of the tag.
X-Forwarded-For (XFF)	■ **Insert** - If the XFF HTTP header is not present in the incoming request, the load balancer inserts a new XFF header with the client IP address. If the XFF HTTP header is present in the incoming request, the load balancer appends the XFF header with the client IP address. ■ **Replace** - If the XFF HTTP header is present in the incoming request, the load balancer replaces the header. Web servers log each request they handle with the requesting client IP address. These logs are used for debugging and analytic purposes. If the deployment topology requires SNAT on the load balancer, then server uses the client SNAT IP address which defeats the purpose of logging.

Option	Description
	As a workaround, the load balancer can be configured to insert XFF HTTP header with the original client IP address. Servers can be configured to log the IP address in the XFF header instead of the source IP address of the connection.
Request Body Size	Enter value for the maximum size of the buffer used to store the HTTP request body. If the size is not specified, then the request body size is unlimited.
Response Timeout (sec)	Enter the time in seconds on how long the load balancer waits for Server HTTP Response before it stops and closes the connection to the pool member and retries the request to another server.
Server Keep-Alive	Toggle the button for the load balancer to turn off TCP multiplexing and enable HTTP keep-alive. If the client uses HTTP/1.0, the load balancer upgrades to HTTP/1.1 protocol and the HTTP keep-alive is set. All HTTP requests received on the same client-side TCP connection are sent to the same server over a single TCP connection to ensure that reauthorization is not required. When HTTP keep-alive is enabled and forwarding rules are configured in the load balancer, the server keep-alive setting takes precedence. As a result, HTTP requests are sent to servers already connected with keep-alive. If you always want to give priority to the forwarding rules when the load balancer rule conditions are met, disable the keep-alive setting. Note that the persistence setting takes precedence over the keep-alive setting. Processing is done in the order of Persistence > Keep-Alive > Load Balancer Rules

d Click **OK**.

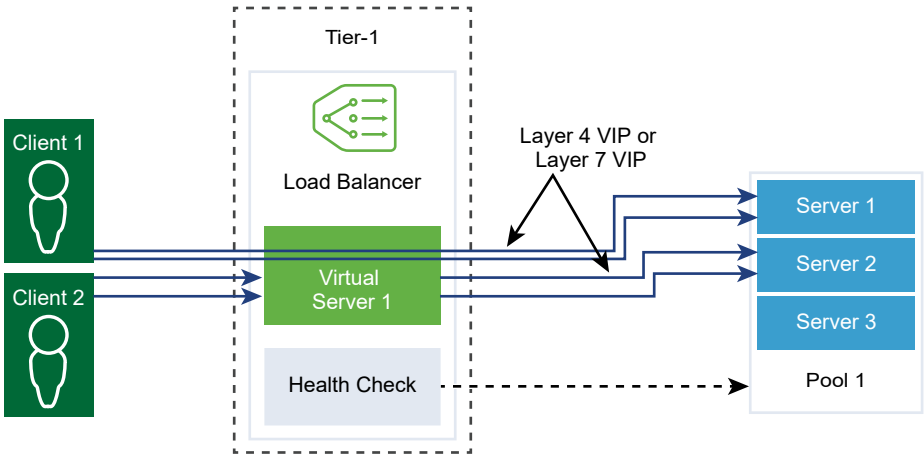
Configure Persistent Profiles in Manager Mode

To ensure stability of stateful applications, load balancers implement persistence which directs all related connections to the same server. Different types of persistence are supported to address different types of application needs.

Some applications maintain the server state such as, shopping carts. Such state might be per client and identified by the client IP address or per HTTP session. Applications might access or modify this state while processing subsequent related connections from the same client or HTTP session.

Source IP persistence profile tracks sessions based on the source IP address. When a client requests a connection to a virtual server that enables the source address persistence, the load balancer checks if that client was previously connected, if so, returns the client to the same server. If not, you can select a server pool member based on the pool load balancing algorithm. Source IP persistence profile is used by Layer 4 and Layer 7 virtual servers.

Cookie persistence profile inserts a unique cookie to identify the session the first time a client accesses the site. The HTTP cookie is forwarded by the client in subsequent requests and the load balancer uses that information to provide the cookie persistence. Cookie persistence profile can only be used by Layer 7 virtual servers. Note that blank space in a cookie name is **not** supported.



Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Networking > Load Balancing > Profiles > Persistence Profiles**.

3 Create a Source IP persistence profile.

   a Select **Add > Source IP Persistence** from the drop-down menu.

   b Enter a name and a description for the Source IP persistence profile.

c Complete the persistence profile details.

You can also accept the default Source IP profile settings.

Option	Description
Share Persistence	Toggle the button to share the persistence so that all virtual servers this profile is associated with can share the persistence table. If persistence sharing is not enabled in the Source IP persistence profile associated to a virtual server, each virtual server that the profile is associated to maintain a private persistence table.
Persistence Entry Timeout	Enter the persistence expiration time in seconds. The load balancer persistence table maintains entries to record that client requests are directed to the same server. ■ If no new connection requests are received from the same client within the timeout period, the persistence entry expires and is deleted. ■ If a new connection request from the same client is received within the timeout period, the timer is reset, and the client request is sent to a sticky pool member. After the timeout period has expired, new connection requests are sent to a server allocated by the load balancing algorithm. For the L7 load balancing TCP source IP persistence scenario, the persistence entry times out if no new TCP connections are made for some time, even if the existing connections are still alive.
HA Persistence Mirroring	Toggle the button to synchronize persistence entries to the HA peer.
Purge Entries When Full	Purge entries when the persistence table is full. A large timeout value might lead to the persistence table quickly filling up when the traffic is heavy. When the persistence table fills up, the oldest entry is deleted to accept the newest entry.

d Click **OK**.

4 Create a Cookie persistence profile.

a Select **Add > Cookie Persistence** from the drop-down menu.

b Enter a name and a description for the Cookie persistence profile.

c Toggle the **Share Persistence** button to share persistence across multiple virtual servers that are associated to the same pool members.

The Cookie persistence profile inserts a cookie with the format, *<name>.<profile-id>.<pool-id>*.

If the persistence shared is not enabled in the Cookie persistence profile associated with a virtual server, the private Cookie persistence for each virtual server is used and is qualified by the pool member. The load balancer inserts a cookie with the format, *<name>.<virtual_server_id>.<pool_id>*.

d Click **Next**.

e Complete the persistence profile details.

Option	Description
Cookie Mode	Select a mode from the drop-down menu. ■ INSERT - Adds a unique cookie to identify the session. ■ PREFIX - Appends to the existing HTTP cookie information. ■ REWRITE - Rewrites the existing HTTP cookie information.
Cookie Name	Enter the cookie name. A blank space in a cookie name is **not** supported.
Cookie Domain	Enter the domain name. HTTP cookie domain can be configured only in the INSERT mode.
Cookie Path	Enter the cookie URL path. HTTP cookie path can be set only in the INSERT mode.
Cookie Garbling	Encrypt the cookie server IP address and port information. Toggle the button to disable encryption. When garbling is disabled, the cookie server IP address and port information is in a plain text.
Cookie Fallback	Select a new server to handle a client request if the cookie points to a server that is in a DISABLED or is in a DOWN state. Toggle the button so that the client request is rejected if cookie points to a server that is in a DISABLED or is in a DOWN state.

f Complete the Cookie expiry details.

Option	Description
Cookie Time Type	Select a cookie time type from the drop-down menu. **Session Cookie** is not stored and will be lost when the browser is closed. **Persistence Cookie** is stored by the browser and is not lost when the browser is closed.
Maximum Idle Time	Enter the time in seconds that a cookie can be idle before it expires.
Maximum Cookie Age	For **Session Cookie** only. Enter the maximum age in seconds that a cookie can be active.

g Click **Finish**.

Configure SSL Profile in Manager Mode

SSL profiles configure application-independent SSL properties such as, cipher lists and reuse these lists across multiple applications. SSL properties are different when the load balancer is acting as a client and as a server, as a result separate SSL profiles for client-side and server-side are supported.

Note SSL profile is not supported in the NSX limited export release.

Client-side SSL profile refers to the load balancer acting as an SSL server and terminating the client SSL connection. Server-side SSL profile refers to the load balancer acting as a client and establishing a connection to the server.

You can specify a cipher list on both the client-side and server-side SSL profiles.

SSL session caching allows the SSL client and server to reuse previously negotiated security parameters avoiding the expensive public key operation during the SSL handshake. SSL session caching is disabled by default on both the client-side and server-side.

SSL session tickets are an alternate mechanism that allow the SSL client and server to reuse previously negotiated session parameters. In SSL session tickets, the client and server negotiate whether they support SSL session tickets during the handshake exchange. If supported by both, server can send an SSL ticket, which includes encrypted SSL session parameters to the client. The client can use that ticket in subsequent connections to reuse the session. SSL session tickets are enabled on the client-side and disabled on the server-side.

Figure 25-17. SSL Offloading

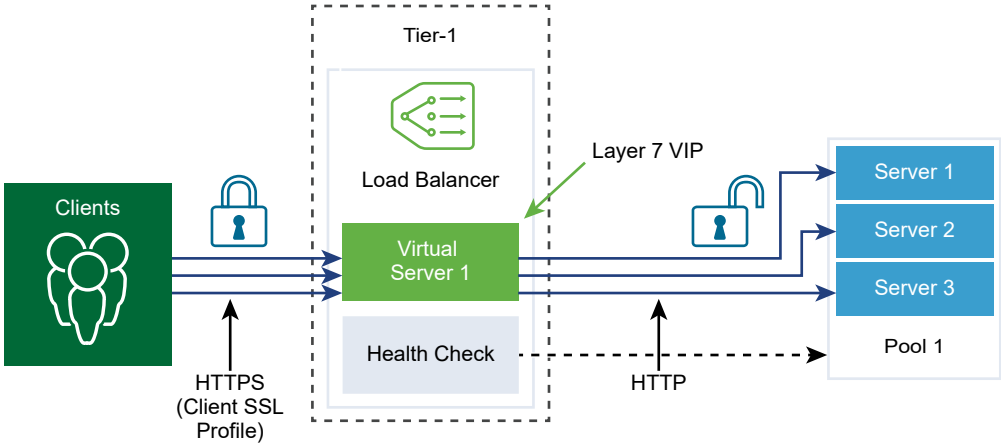
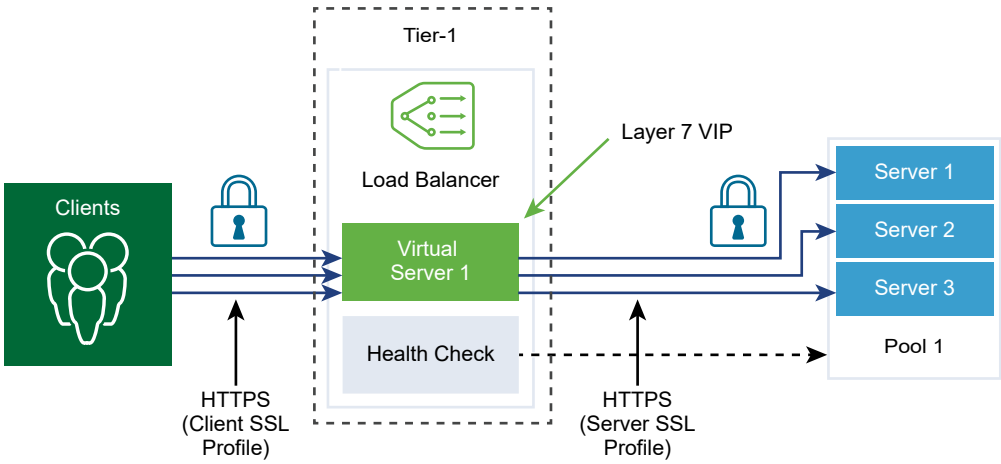


Figure 25-18. End-to-End SSL



Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Networking > Load Balancing > Profiles > SSL Profiles**.

3 Create a Client SSL profile.

a Select **Add > Client Side SSL** from the drop-down menu.

b Enter a name and a description for the Client SSL profile.

c Assign the SSL Ciphers to be included in the Client SSL profile.

You can also create custom SSL Ciphers.

d Click the arrow to move the ciphers to the Selected section.

e Click the **Protocols and Sessions** tab.

f Select the SSL protocols to be included in the Client SSL profile.

SSL protocol versions TLS1.1 and TLS1.2 are enabled by default. TLS1.0 is also supported, but disabled by default.

g Click the arrow to move the protocol to the Selected section.

h Complete the SSL protocol details.

You can also accept the default SSL profile settings.

Option	Description
Session Caching	SSL session caching allows the SSL client and server to reuse previously negotiated security parameters avoiding the expensive public key operation during an SSL handshake.
Session Cache Entry Timeout	Enter the cache timeout in seconds to specify how long the SSL session parameters must be kept and can be reused.
Prefer Server Cipher	Toggle the button so that the server can select the first supported cipher from the list it can support. During an SSL handshake, the client sends an ordered list of supported ciphers to the server.

i Click **OK**.

4 Create a Server SSL profile.

a Select **Add > Server Side SSL** from the drop-down menu.

b Enter a name and a description for the Server SSL profile.

c Select the SSL Ciphers to be included in the Server SSL profile.

You can also create custom SSL Ciphers.

d Click the arrow to move the ciphers to the Selected section.

e Click the **Protocols and Sessions** tab.

f Select the SSL protocols to be included in the Server SSL profile.

SSL protocol versions TLS1.1 and TLS1.2 are enabled by default. TLS1.0 is also supported, but disabled by default.

g Click the arrow to move the protocol to the Selected section.

h Accept the default session caching setting.

SSL session caching allows the SSL client and server to reuse previously negotiated security parameters avoiding the expensive public key operation during an SSL handshake.

i Click **OK**.

Configure Layer 4 Virtual Servers in Manager Mode

Virtual servers receive all the client connections and distribute them among the servers. A virtual server has an IP address, a port, and a protocol. For Layer 4 virtual servers, lists of ports ranges can be specified instead of a single TCP or UDP port to support complex protocols with dynamic ports.

A Layer 4 virtual server must be associated to a primary server pool, also called a default pool.

If a virtual server status is disabled, any new connection attempts to the virtual server are rejected by sending either a TCP RST for the TCP connection or ICMP error message for UDP. New connections are rejected even if there are matching persistence entries for them. Active connections continue to be processed. If a virtual server is deleted or disassociated from a load balancer, then active connections to that virtual server fail.

Prerequisites

- Verify that application profiles are available. See Configure Application Profiles in Manager Mode.
- Verify that persistent profiles are available. See Configure Persistent Profiles in Manager Mode.
- Verify that SSL profiles for the client and server are available. See Configure SSL Profile in Manager Mode.
- Verify that server pools are available. See Add a Server Pool for Load Balancing in Manager Mode.
- Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Networking > Load Balancing > Virtual Servers > Add**.

3 Enter a name and a description for the Layer 4 virtual server.

4 Select a Layer 4 protocol from the drop-down menu.

   Layer 4 virtual servers support either the Fast TCP or Fast UDP protocol, but not both. For Fast TCP or Fast UDP protocol support on the same IP address and port, for example DNS, a virtual server must be created for each protocol.

   Based on the protocol type, the existing application profile is automatically populated.

5 Toggle the Access Log button to enable logging for the Layer 4 virtual server.

6 Click **Next**.

7 Enter the virtual server IP address and port number.

   You can enter the virtual server port number or port range.

8 Complete the advanced properties details.

Option	Description
Maximum Concurrent Connection	Set the maximum concurrent connection allowed to a virtual server so that the virtual server does not deplete resources of other applications hosted on the same load balancer.
Maximum New Connection Rate	Set the maximum new connection to a server pool member so that a virtual server does not deplete resources.
Default Pool Member Port	Enter a default pool member port if the pool member port for a virtual server is not defined. For example, if a virtual server is defined with port range 2000-2999 and the default pool member port range is set as 8000-8999, then an incoming client connection to the virtual server port 2500 is sent to a pool member with a destination port set to 8500.

9 Select an existing server pool from the drop-down menu.

   The server pool consists of one or more servers, also called pool members that are similarly configured and running the same application.

10 Select an existing sorry server pool from the drop-down menu.

   The sorry server pool serves the request when a load balancer cannot select a backend server to the serve the request from the default pool.

11 Click **Next**.

12 Select the existing persistence profile from the drop-down menu.

   Persistence profile can be enabled on a virtual server to allow related client connections to be sent to the same server.

13 Click **Finish**.

Configure Layer 7 Virtual Servers in Manager Mode

Virtual servers receive all the client connections and distribute them among the servers. A virtual server has an IP address, a port, and a protocol TCP.

Load balancer rules are supported for only Layer 7 virtual servers with an HTTP application profile. Different load balancer services can use load balancer rules.

Each load balancer rule consists of single or multiple match conditions and single or multiple actions. If the match conditions are not specified, then the load balancer rule always matches and is used to define default rules. If more than one match condition is specified, then the matching strategy determines if all conditions must match or any one condition must match for the load balancer rule to be considered a match.

Each load balancer rule is implemented at a specific phase of the load balancing processing; HTTP Request Rewrite, HTTP Request Forwarding, and HTTP Response Rewrite. Not all the match conditions and actions are applicable to each phase.

If a virtual server status is disabled, any new connection attempts to the virtual server are rejected by sending either a TCP RST for the TCP connection or ICMP error message for UDP. New connections are rejected even if there are matching persistence entries for them. Active connections continue to be processed. If a virtual server is deleted or disassociated from a load balancer, then active connections to that virtual server fail.

Prerequisites

- Verify that application profiles are available. See Configure Application Profiles in Manager Mode.
- Verify that persistent profiles are available. See Configure Persistent Profiles in Manager Mode.
- Verify that SSL profiles for the client and server are available. See Configure SSL Profile in Manager Mode.
- Verify that server pools are available. See Add a Server Pool for Load Balancing in Manager Mode.
- Verify that CA and client certificate are available. See Create a Certificate Signing Request File.
- Verify that a certification revocation list (CRL) is available. See Import a Certificate Revocation List.
- Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.
- Configure Layer 7 Virtual Server Pool and Rules

  With Layer 7 virtual servers, you can optionally configure load balancer rules and customize load balancing behavior using match or action rules.
- Configure Layer 7 Virtual Server Load Balancing Profiles

  With Layer 7 virtual servers, you can optionally configure load balancer persistence, client-side SSL, and server-side SSL profiles.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **Networking > Load Balancing > Virtual Servers > Add**.

3 Enter a name and a description for the Layer 7 virtual server.

4 Select the Layer 7 menu item.

Layer 7 virtual servers support the HTTP and HTTPS protocols.

The existing HTTP application profile is automatically populated.

5 (Optional) Click **Next** to configure server pool and load balancing profiles.

6 Click **Finish**.

Configure Layer 7 Virtual Server Pool and Rules

With Layer 7 virtual servers, you can optionally configure load balancer rules and customize load balancing behavior using match or action rules.

Load Balancer rules support REGEX for match types. PCRE style REGEX patters is supported with a few limitations on advanced use cases. When REGEX is used in match conditions, named capturing groups are supported.

REGEX restrictions include:

- Character unions and intersections are not supported. For example, do not use [a-z[0-9]] and [a-z&&[aeiou]] instead use [a-z0-9] and [aeiou] respectively.
- Only 9 back references are supported and \1 through \9 can be used to refer to them.
- Use \0dd format to match octal characters, not the \ddd format.
- Embedded flags are not supported at the top level, they are only supported within groups. For example, do not use "Case (?i:s)ensitive" instead use "Case ((?i:s)ensitive)".
- Preprocessing operations \l, \u, \L, \U are not supported. Where \l - lowercase next char \u - uppercase next char \L - lower case until \E \U - upper case to \E.
- (?(condition)X), (?{code}), (??{Code}) and (?#comment) are not supported.
- Predefined Unicode character class \X is not supported
- Using named character construct for Unicode characters is not supported. For example, do not use \N{name} instead use \u2018.

When REGEX is used in match conditions, named capturing groups are supported. For example, REGEX match pattern /news/(?<year>\d+)-(?<month>\d+)-(?<day>\d+)/(?<article>.*) can be used to match a URI like /news/2018-06-15/news1234.html.

Then variables are set as follows, $year = "2018" $month = "06" $day = "15" $article = "news1234.html". After the variables are set, these variables can be used in load balancer rule actions. For example, URI can be rewritten using the matched variables like, /news.py?year=$year&month=$month&day=$day&article=$article. Then the URI gets rewritten as /news.py?year=2018&month=06&day=15&article=news1234.html.

Rewrite actions can use a combination of named capturing groups and built-in variables. For example, URI can be written as /news.py?year=$year&month=$month&day=$day&article=$article&user_ip=$_remote_addr. Then the example URI gets rewritten as /news.py?year=2018&month=06&day=15&article=news1234.html&user_ip=1.1.1.1.

Note For named capturing groups, the name cannot start with an _ character.

In addition to named capturing groups, the following built-in variables can be used in rewrite actions. All the built-in variable names start with _.

- $_args - arguments from the request
- $_arg_<name> - argument <name> in the request line
- $_cookie_<name> - value of <name> cookie
- $_upstream_cookie_<name> - cookie with the specified name sent by the upstream server in the "Set-Cookie" response header field
- $_upstream_http_<name> - arbitrary response header field and <name> is the field name converted to lower case with dashes replaced by underscores
- $_host - in the order of precedence - host name from the request line, or host name from the "Host" request header field, or the server name matching a request
- $_http_<name> - arbitrary request header field and <name> is the field name converted to lower case with dashes replaced by underscores
- $_https - "on" if connection operates in SSL mode, or "" otherwise
- $_is_args - "?" if a request line has arguments, or "" otherwise
- $_query_string - same as $_args
- $_remote_addr - client address
- $_remote_port - client port
- $_request_uri - full original request URI (with arguments)
- $_scheme - request scheme, "http" or "https"
- $_server_addr - address of the server which accepted a request
- $_server_name - name of the server which accepted a request
- $_server_port - port of the server which accepted a request
- $_server_protocol - request protocol, usually "HTTP/1.0" or "HTTP/1.1"
- $_ssl_client_cert - returns the client certificate in the PEM format for an established SSL connection, with each line except the first prepended with the tab character
- $_ssl_server_name - returns the server name requested through SNI
- $_uri - URI path in request

- $_ssl_ciphers: returns the client SSL ciphers
- $_ssl_client_i_dn: returns the "issuer DN" string of the client certificate for an established SSL connection according to RFC 2253
- $_ssl_client_s_dn: returns the "subject DN" string of the client certificate for an established SSL connection according to RFC 2253
- $_ssl_protocol: returns the protocol of an established SSL connection
- $_ssl_session_reused: returns "r" if an SSL session was reused, or "." otherwise

Prerequisites

Verify a Layer 7 virtual server is available. See Configure Layer 7 Virtual Servers in Manager Mode.

Procedure

1 Open the Layer 7 virtual server.

2 Skip to the Virtual Server Identifiers page.

3 Enter the virtual server IP address and port number.

You can enter the virtual server port number or port range.

4 Complete the advanced properties details.

Option	Description
Maximum Concurrent Connection	Set the maximum concurrent connection allowed to a virtual server so that the virtual server does not deplete resources of other applications hosted on the same load balancer.
Maximum New Connection Rate	Set the maximum new connection to a server pool member so that a virtual server does not deplete resources.
Default Pool Member Port	Enter a default pool member port if the pool member port for a virtual server is not defined. For example, if a virtual server is defined with port range 2000–2999 and the default pool member port range is set as 8000-8999, then an incoming client connection to the virtual server port 2500 is sent to a pool member with a destination port set to 8500.

5 (Optional) Select an existing default server pool from the drop-down menu.

The server pool consists of one or more servers, called pool members that are similarly configured and running the same application.

6 Click **Add** to configure the load balancer rules for the HTTP Request Rewrite phase.

Supported match types are, REGEX, STARTS_WITH, ENDS_WITH, etc and inverse option.

Supported Match Condition	Description
HTTP Request Method	Match an HTTP request method. http_request.method - value to match
HTTP Request URI	Match an HTTP request URI without query arguments. http_request.uri - value to match
HTTP Request URI arguments	Match an HTTP request URI query argument. http_request.uri_arguments - value to match
HTTP Request Version	Match an HTTP request version. http_request.version - value to match
HTTP Request Header	Match any HTTP request header. http_request.header_name - header name to match http_request.header_value - value to match
HTTP Request Payload	Match an HTTP request body content. http_request.body_value - value to match
TCP Header Fields	Match a TCP source or the destination port. tcp_header.source_port - source port to match tcp_header.destination_port - destination port to match
IP Header Fields	Match an IP source or destination address. ip_header.source_address - source address to match ip_header.destination_address - destination address to match

Action	Description
HTTP Request URI Rewrite	Modify an URI. http_request.uri - URI (without query arguments) to write http_request.uri_args - URI query arguments to write
HTTP Request Header Rewrite	Modify value of an HTTP header. http_request.header_name - header name http_request.header_value - value to write

7 Click **Add** to configure the load balancer rules for the HTTP Request Forwarding.

All match values accept regular expressions.

Supported Match Condition	Description
HTTP Request Method	Match an HTTP request method. http_request.method - value to match
HTTP Request URI	Match an HTTP request URI. http_request.uri - value to match
HTTP Request URI args	Match an HTTP request URI query argument. http_request.uri_args - value to match

Supported Match Condition	Description
HTTP Request Version	Match an HTTP request version. http_request.version - value to match
HTTP Request Header	Match any HTTP request header. http_request.header_name - header name to match http_request.header_value - value to match
HTTP Request Payload	Match an HTTP request body content. http_request.body_value - value to match
TCP Header Fields	Match a TCP source or the destination port. tcp_header.source_port - source port to match tcp_header.destination_port - destination port to match
IP Header Fields	Match an IP source address. ip_header.source_address - source address to match

Action	Description
Reject	Reject a request, for example, by setting status to 5xx. http_forward.reply_status - HTTP status code used to reject http_forward.reply_message - HTTP rejection message
Redirect	Redirect a request. Status code must be set to 3xx. http_forward.redirect_status - HTTP status code for redirect http_forward.redirect_url - HTTP redirect URL
Select Pool	Force the request to a specific server pool. Specified pool member's configured algorithm (predictor) is used to select a server within the server pool. http_forward.select_pool - server pool UUID

8 Click **Add** to configure the load balancer rules for the HTTP Response Rewrite.

All match values accept regular expressions.

Supported Match Condition	Description
HTTP Response Header	Match any HTTP response header. http_response.header_name - header name to match http_response.header_value - value to match

Action	Description
HTTP Response Header Rewrite	Modify the value of an HTTP response header. http_response.header_name - header name http_response.header_value - value to write

9 (Optional) Click **Next** to configure load balancing profiles.

10 Click **Finish**.

Configure Layer 7 Virtual Server Load Balancing Profiles

With Layer 7 virtual servers, you can optionally configure load balancer persistence, client-side SSL, and server-side SSL profiles.

Note SSL profile is not supported in the NSX limited export release.

If a client-side SSL profile binding is configured on a virtual server but not a server-side SSL profile binding, then the virtual server operates in an SSL-terminate mode, which has an encrypted connection to the client and plain text connection to the server. If both the client-side and server-side SSL profile bindings are configured, then the virtual server operates in SSL-proxy mode, which has an encrypted connection both to the client and the server.

Associating server-side SSL profile binding without associating a client-side SSL profile binding is currently not supported. If a client-side and a server-side SSL profile binding is not associated with a virtual server and the application is SSL-based, then the virtual server operates in an SSL-unaware mode. In this case, the virtual server must be configured for Layer 4. For example, the virtual server can be associated to a fast TCP profile.

Prerequisites

Verify a Layer 7 virtual server is available. See Configure Layer 7 Virtual Servers in Manager Mode.

Procedure

1. Open the Layer 7 virtual server.
2. Skip to the Load Balancing Profiles page.
3. Toggle the Persistence button to enable the profile.

   Persistence profile allows related client connections to be sent to the same server.
4. Select either the Source IP Persistence or Cookie Persistence profile.
5. Select the existing persistence profile from the drop-down menu.
6. Click **Next**.
7. Toggle the Client Side SSL button to enable the profile.

   Client-side SSL profile binding allows multiple certificates, for different host names to be associated to the same virtual server.

   The associated Client-side SSL profile is automatically populated.
8. Select a default certificate from the drop-down menu.

   This certificate is used if the server does not host multiple host names on the same IP address or if the client does not support Server Name Indication (SNI) extension.
9. Select the available SNI certificate and click the arrow to move the certificate to the Selected section.
10. (Optional) Toggle the Mandatory Client Authentication to enable this menu item.

11 Select the available CA certificate and click the arrow to move the certificate to the Selected section.

12 Set the certificate chain depth to verify the depth in the server certificates chain.

13 Select the available CRL and click the arrow to move the certificate to the Selected section.

A CRL can be configured to disallow compromised server certificates.

14 Click **Next**.

15 Toggle the Server Side SSL button to enable the profile.

The associated Server-side SSL profile is automatically populated.

16 Select a client certificate from the drop-down menu.

The client certificate is used if the server does not host multiple host names on the same IP address or if the client does not support Server Name Indication (SNI) extension.

17 Select the available SNI certificate and click the arrow to move the certificate to the Selected section.

18 (Optional) Toggle the Server Authentication to enable this menu item.

Server-side SSL profile binding specifies whether the server certificate presented to the load balancer during the SSL handshake must be validated or not. When validation is enabled, the server certificate must be signed by one of the trusted CAs whose self-signed certificates are specified in the same server-side SSL profile binding.

19 Select the available CA certificate and click the arrow to move the certificate to the Selected section.

20 Set the certificate chain depth to verify the depth in the server certificates chain.

21 Select the available CRL and click the arrow to move the certificate to the Selected section.

A CRL can be configured to disallow compromised server certificates. OCSP and OCSP stapling are not supported on the server-side.

22 Click **Finish**.

Firewall in Manager Mode

You can configure Distributed Firewall and logical router Firewall in **Manager Mode**.

Note If you use **Manager** mode to modify objects created in the **Policy** mode, some settings might not be configurable. These read-only settings have this icon next to them: ⊖. See Chapter 1 NSX Manager for more information.

Add or Delete a Firewall Rule to a Logical Router in Manager Mode

You can add firewall rules to a tier-0 or tier-1 logical router to control communication into the router.

Edge fire-walling is implemented on uplink router ports, meaning that firewall rules will be applicable only if traffic hits uplink router ports on edge. To apply firewall rules to particular IP destination, you must configure groups with /32 network. If you provide a subnet other than /32, firewall rules will be applied to the complete subnet.

Prerequisites

- Familiarize yourself with the parameters of a firewall rule. See Add a Firewall Rule in Manager Mode.
- Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Locate the router in **Networking > Tier-0 Logical Routers** or **Networking > Tier-1 Logical Routers**.
3. Click the name of the logical router.
4. Select **Services > Edge Firewall**.
5. Click an existing section or rule.
6. To add a rule, click **Add Rule** on the menu bar and select **Add Rule Above** or **Add Rule Below**, or click the menu icon in the first column of a rule and select **Add Rule Above** or **Add Rule Below**, and specify the rule parameters.

   The Applied To field is not shown because this rule applies only to the logical router.
7. To delete a rule, select the rule, click **Delete** on the menu bar or click the menu icon in the first column and select **Delete**.

Results

Note If you add a firewall rule to a tier-0 logical router and the NSX Edge cluster backing the router is running in active-active mode, the firewall can only run in stateless mode. If you configure the firewall rule with stateful services such as HTTP, SSL, TCP, and so on, the firewall rule will not work as expected. To avoid this issue, configure the NSX Edge cluster to run in active-standby mode.

Configure Firewall for a Logical Switch Bridge Port in Manager Mode

You can configure firewall sections and firewall rules for the bridge port of a layer 2 bridge-backed logical switch. The bridge must be created using NSX Edge nodes.

Prerequisites

- Verify that the switch is attached to a bridge profile. See Extend an Overlay Segment to a VLAN or a Range of VLANs in Manager Mode.

- Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Select **Security > Bridge Firewall**.
3. Select a logical switch.

   The switch must be attached to a bridge profile.
4. Follow the same steps in previous sections for configuring layer 2 or layer 3 firewall.

Firewall Sections and Firewall Rules

Firewall sections are used to group a set of firewall rules.

A firewall section is made up from one or more individual firewall rules. Each individual firewall rule contains instructions that determine whether a packet should be allowed or blocked; which protocols it is allowed to use; which ports it is allowed to use and so forth. Sections are used for multi-tenancy , such as specific rules for sales and engineering departments in separate sections.

A section can be defined as enforcing stateful or stateless rules. Stateless rules are treated as traditional stateless ACLs. Reflexive ACLs are not supported for stateless sections. A mix of stateless and stateful rules on a single logical switch port is not recommended and may cause undefined behavior.

Rules can be moved up and down within a section. For any traffic attempting to pass through the firewall, the packet information is subjected to the rules in the order shown in the section, beginning at the top and proceeding to the default rule at the bottom. The first rule that matches the packet has its configured action applied, and any processing specified in the rule's configured options is performed and all subsequent rules are ignored (even if a later rule is a better match). Thus, you should place specific rules above more general rules to ensure those rules are not ignored. The default rule, located at the bottom of the rule table, is a "catchall" rule; packets not matching any other rules will be enforced by the default rule.

Activate and Deactivate Distributed Firewall in Manager Mode

You can activate and deactivate the distributed firewall feature.

If it is deactivated, no firewall rules are enforced at the dataplane level. Upon reactivation rules are enforced.

Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1 Navigate to **Security > Distributed Firewall**.

2 Click the **Settings** tab.

3 Click Distributed Firewall **Edit**.

4 In the dialog box, toggle the firewall status to green (activated) or gray (deactivated).

5 Click **Save**.

Add a Firewall Rule Section in Manager Mode

A firewall rule section is edited and saved independently and is used to apply separate firewall configuration to tenants.

Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1 Select **Security > Distributed Firewall**.

2 Click the **General** tab for layer 3 (L3) rules or the **Ethernet** tab for layer 2 (L2) rules.

3 Click an existing section or rule.

4 Click the section icon on the menu bar and select **Add Section Above** or **Add Section Below**.

> **Note** For any traffic attempting to pass through the firewall, the packet information is subjected to the rules in the order shown in the Rules table, beginning at the top and proceeding to the default rules at the bottom. In some cases, the order of precedence of two or more rules might be important in determining the disposition of a packet.

5 Enter the section name.

6 To make the firewall stateless, select the **Enable Stateless Firewall**. This option is applicable for L3 only.

Stateless firewalls watch network traffic, and restrict or block packets based on source and destination addresses or other static values. For TCP and UDP flows, after the first packet, a cache is created and maintained for the traffic tuple in either direction, if the firewall result is ALLOW. This means that the traffic no longer needs to check with the firewall rules, resulting in lower latency. Stateless firewalls are thus typically faster and perform better under heavier traffic loads.

Stateful firewalls can watch traffic streams from end to end. The firewall is always consulted for every packet, to validate the state and sequence numbers. Stateful firewalls are better at identifying unauthorized and forged communications.

There is no toggling between stateful and stateless once it is defined.

7 Select one or more objects to apply the section.

The types of object are logical ports, logical switches, and NSGroups. If you select an NSGroup, it must contain one or more logical switches or logical ports. If the NSGroup contains only IP sets or MAC sets, it will be ignored.

Note If both the section and the rules within have **Applied To** set to NSGroup, then the **Applied To** in a section it will override any **Applied To** settings in the rules in that section. This is because the firewall section level **Applied To** takes precedence over **Applied To** at the rule level.

8 Click **OK**.

What to do next

Add Firewall rules to the section.

Delete a Firewall Rule Section in Manager Mode

A firewall rule section can be deleted when it is no longer used.

When you delete a firewall rule section, all rules in that section are deleted. You cannot delete a section and add it again at a different place in the firewall table. To do so, you must delete the section and publish the configuration. Then add the deleted section to the firewall table and re-publish the configuration.

Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1 Select **Security > Distributed Firewall**.

2 Click the **General** tab for L3 rules or the **Ethernet** tab for L2 rules.

3 Click the menu icon in the first column of the section and select **Delete Section**.

You can also select the section and click the delete icon on the menu bar.

Enable and Disable Section Rules in Manager Mode

You can enable or disable all rules in a firewall rule section.

Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1 Select **Security > Distributed Firewall**.

2 Click the **General** tab for L3 rules or the **Ethernet** tab for L2 rules.

3 Click the menu icon in the first column of the section and select **Enable All Rules** or **Disable All Rules**.

4 Click **Publish**.

Enable and Disable Section Logs in Manager Mode

Enabling logs for section rules records information on packets for all of the rules in a section. Depending on the number of rules in a section, a typical firewall section will generate large amounts of log information and can affect performance.

Logs are stored in the /var/log/dfwpktlogs.log file on ESXi hosts.

Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1 Select **Security > Distributed Firewall**.

2 Click the **General** tab for L3 rules or the **Ethernet** tab for L2 rules.

3 Click the menu icon in the first column of the section and select **Enable Logs** or **Disable Logs**.

4 Click **Publish**.

Configure a Firewall Exclusion List in Manager Mode

A logical port, logical switch, or NSGroup can be excluded from a firewall rule.

After you've created a section with firewall rules you may want to exclude an NSX appliance port from the firewall rules.

Note NSX automatically adds NSX Manager and NSX Edge node virtual machines to the firewall exclusion list.

Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1 Select **Security > Distributed Firewall > Exclusion List > Add**.

2 Select a type and an object.

   The available types are **Logical Port**, **Logical Switch**, and **NSGroup**.

3 Click **OK**.

4 To remove an object from the exclusion list, select the object and click **Delete** on the menu bar.

About Firewall Rules

NSX uses firewall rules to specify traffic handling in and out of the network.

Firewall offers multiple sets of configurable rules: Layer 3 rules (General tab) and Layer 2 rules (Ethernet tab). Layer 2 firewall rules are processed before Layer 3 rules and if allowed in the Layer 2 rules will then be processed by the Layer 3 rules. You can configure an exclusion list that contains logical switches, logical ports, or groups that are to be excluded from firewall enforcement.

Firewall Rules are enforced as follows:

- Rules are processed in top-to-bottom ordering.
- Each packet is checked against the top rule in the rule table before moving down the subsequent rules in the table.
- The first rule in the table that matches the traffic parameters is enforced.

No subsequent rules can be enforced as the search is then terminated for that packet. Because of this behavior, it is always recommended to put the most granular policies at the top of the rule table. This will ensure they will be enforced before more specific rules.

The default rule, located at the bottom of the rule table, is a catchall rule; packets not matching any other rules will be enforced by the default rule. After the host preparation operation, the default rule is set to allow action. This ensures that VM-to-VM communication is not broken during staging or migration phases. It is a best practice to then change this default rule to block action and enforce access control through a positive control model (i.e., only traffic defined in the firewall rule is allowed onto the network).

Note TCP strict can be enabled on a per section basis to turn off mid-session pick-up and enforce the requirement for a three-way handshake. When enabling TCP strict mode for a particular Distributed Firewall Section, and using a default ANY-ANY Block rule, packets that do not complete the three-way handshake connection requirements, and that match a TCP-based rule in this section are dropped. Strict is only applied to stateful TCP rules and is enabled at the distributed firewall section level. TCP strict is not enforced for packets that match a default ANY-ANY Allow which as no TCP service specified.

Table 25-1. Properties of a Firewall Rule

Property	Description
Name	Name of the firewall rule.
ID	Unique system generated ID for each rule.
Source	The source of the rule can be either an IP or MAC address or an object other than an IP address. The source will match any if not defined. Both IPv4 and IPv6 are supported for source or destination range.
Destination	The destination IP or MAC address/netmask of the connection that is affected by the rule. The destination will match any if not defined. Both IPv4 and IPv6 are supported for source or destination range.
Service	The service can be a predefined port protocol combination for L3. For L2 it can be ether-type. For both L2 and L3 you can manually define a new service or service group. The service will match any, if it is not specified.
Applied To	Defines the scope at which this rule is applicable. If not defined the scope will be all logical ports. If you have added "applied to" in a section it will overwrite the rule.
Log	Logging can be turned off or on. Logs are stored at /var/log/dfwpktlogs.log file on ESXi hosts.
Action	The action applied by the rule can be **`Allow`**, **`Drop`**, or **`Reject`**. The default is **`Allow`**.
IP Protocol	The options are **`IPv4`**, **`IPv6`**, and **`IPv4_IPv6`**. The default is **`IPv4_IPv6`**. To access this property, click the **Advanced Settings** icon.
Direction	The options are **`In`**, **`Out`**, and **`In/Out`**. The default is **`In/Out`**. This field refers to the direction of traffic from the point of view of the destination object. **`In`** means that only traffic to the object is checked, **`Out`** means that only traffic from the object is checked, and **`In/Out`** means traffic in both directions is checked. To access this property, click the **Advanced Settings** icon.
Rule Tags	Tags that have been added to the rule. To access this property, click the **Advanced Settings** icon.
Flow Statistics	Read-only field that displays the byte, packet count, and sessions. To access this property, click the graph icon.

Note If SpoofGuard is not enabled, automatically discovered address bindings cannot be guaranteed to be trustworthy because a malicious virtual machine can claim the address of another virtual machine. SpoofGuard, if enabled, verifies each discovered binding so that only approved bindings are presented.

Add a Firewall Rule in Manager Mode

A firewall is a network security system that monitors and controls the incoming and outgoing network traffic based on predetermined firewall rules.

Firewall rules are added at the NSX Manager scope. Using the Applied To field, you can then narrow down the scope at which you want to apply the rule. You can add multiple objects at the source and destination levels for each rule, which helps reduce the total number of firewall rules to be added.

Note By default, a rule matches on the default of any source, destination, and service rule elements, matching all interfaces and traffic directions. If you want to restrict the effect of the rule to particular interfaces or traffic directions, you must specify the restriction in the rule.

Prerequisites

- To use a group of addresses, first manually associate the IP and MAC address of each VM with their logical switch.
- Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1 Select **Security > Distributed Firewall**.

2 Click the **General** tab for L3 rules or the **Ethernet** tab for L2 rules.

3 Click an existing section or rule.

4 Click the menu icon in the first column of a rule and select **Add Rule Above** or **Add Rule Below**.

A new row appears to define a firewall rule.

Note For any traffic attempting to pass through the firewall, the packet information is subjected to the rules in the order shown in the Rules table, beginning at the top and proceeding to the default rules at the bottom. In some cases, the order of precedence of two or more rules might be important in determining the disposition of a packet.

5 In the **Name** column, enter the rule name.

6 In the **Source** column, click the edit icon and select the source of the rule. The source will match any if not defined.

Option	Description
IP Addresses	Enter multiple IP or MAC addresses in a comma-separated list. The list can contain up to 255 characters. Both IPv4 and IPv6 formats are supported.
Container Objects	The available objects are IP Set, Logical Port, Logical Switch, and NS Group. Select the objects and click **OK**.

7 In the **Destination** column, click the edit icon and select the destination. The destination will match any if not defined.

Option	Description
IP Addresses	You can enter multiple IP or MAC addresses in a comma-separated list. The list can contain up to 255 characters. Both IPv4 and IPv6 formats are supported.
Container Objects	The available objects are IP Set, Logical Port, Logical Switch, and NS Group. Select the objects and click **OK**.

8 In the **Service** column, click the edit icon and select services. The service will match any if not defined.

9 To select a predefined service, select one of more available services.

10 To define a new service, click the **Raw Port-Protocol** tab and click **Add**..

Option	Description
Type of Service	■ ALG ■ ICMP ■ IGMP ■ IP ■ L4 Port Set
Protocol	Select one of the available protocols.
Source Ports	Enter the source port.
Destination Ports	Select the destination port.

11 In the **Applied To** column, click the edit icon and select objects.

12 In the **Log** column, set the logging option.

Logs are in the `/var/log/dfwpktlogs.log` file on ESXi. Enabling logging can affect performance.

13 In the **Action** column, select an action.

Option	Description
Allow	Allows all L3 or L2 traffic with the specified source, destination, and protocol to pass through the current firewall context. Packets that match the rule, and are accepted, traverse the system as if the firewall is not present
Drop	Drops packets with the specified source, destination, and protocol. Dropping a packet is a silent action with no notification to the source or destination systems. Dropping the packet causes the connection to be retried until the retry threshold is reached.
Reject	Rejects packets with the specified source, destination, and protocol. Rejecting a packet is a more graceful way to deny a packet, as it sends a destination unreachable message to the sender. If the protocol is TCP, a TCP RST message is sent. ICMP messages with administratively prohibited code are sent for UDP, ICMP, and other IP connections. One benefit of using Reject is that the sending application is notified after only one attempt that the connection cannot be established.

14 Click the **Advanced Settings** icon to specify IP protocol, direction, rule tags, and comments.

15 Click **Publish**.

Delete a Firewall Rule in Manager Mode

A firewall is a network security system that monitors and controls the incoming and outgoing network traffic based on predetermined firewall rules. Custom defined rules can be added and deleted.

Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1 Select **Security > Distributed Firewall**.

2 Click the **General** tab for L3 rules or the **Ethernet** tab for L2 rules.

3 Click the menu icon in the first column of the rule and select **Delete Rule**.

4 Click **Publish**.

Edit the Default Distributed Firewall Rule in Manager Mode

You can edit the default firewall settings that apply to traffic that does not match any of the user-defined firewall rules.

The default firewall rules apply to traffic that does not match any of the user-defined firewall rules. The default Layer 3 rule is under the **General** tab and the default Layer 2 rule is under the **Ethernet** tab.

The default firewall rules allow all L3 and L2 traffic to pass through all prepared clusters in your infrastructure. The default rule is always at the bottom of the rules table and cannot be deleted. However, you can change the **Action** element of the rule from **Allow** to **Drop** or **Reject**, and indicate whether traffic for that rule should be logged.

The default Layer 3 firewall rule applies to all traffic, including DHCP. If you change the **Action** to **Drop** or **Reject**, DHCP traffic will be blocked. You will need to create a rule to allow DHCP traffic.

Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1. Select **Security > Distributed Firewall**.
2. Click the **General** tab for L3 rules or the **Ethernet** tab for L2 rules.
3. In the **Name** column, enter a new name.
4. In the **Action** column, select one of the options.
   - Allow - Allows all L3 or L2 traffic with the specified source, destination, and protocol to pass through the current firewall context. Packets that match the rule, and are accepted, traverse the system as if the firewall is not present.
   - Drop - Drops packets with the specified source, destination, and protocol. Dropping a packet is a silent action with no notification to the source or destination systems. Dropping the packet causes the connection to be retried until the retry threshold is reached.
   - Reject - Rejects packets with the specified source, destination, and protocol. Rejecting a packet is a more graceful way to deny a packet, as it sends a destination unreachable message to the sender. If the protocol is TCP, a TCP RST message is sent. ICMP messages with administratively prohibited code are sent for UDP, ICMP, and other IP connections. One benefit of using Reject is that the sending application is notified after only one attempt that the connection cannot be established.
5. In the **Log**, enable or disable logging.

   Enabling logging can affect performance.
6. Click **Publish**.

Change the Order of a Firewall Rule in Manager Mode

Rules are processed in top-to-bottom ordering. You can change the order of the rules in the list.

For any traffic attempting to pass through the firewall, the packet information is subjected to the rules in the order shown in the rules table, beginning at the top and proceeding to the default rules at the bottom. In some cases, the order of precedence of two or more rules might be important in determining the traffic flow.

You can move a custom rule up or down in the table. The default rule is always at the bottom of the table and cannot be moved.

Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1. Select **Security > Distributed Firewall**.
2. Click the **General** tab for L3 rules or the **Ethernet** tab for L2 rules.
3. Select the rule and click the **Move Up** or **Move Down** icon on the menu bar.
4. Click **Publish**.

Filter Firewall Rules in Manager Mode

When you navigate to the firewall section, initially all the rules are displayed. You can apply a filter to control what is displayed so that you see only a subset of the rules. This can make it easier to manage the rules.

Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See Chapter 1 NSX Manager. If you do not see the **Policy** and **Manager** mode buttons, see Configure the User Interface Settings.

Procedure

1. Select **Security > Distributed Firewall**.
2. Click the **General** tab for L3 rules or the **Ethernet** tab for L2 rules.
3. In the search text field on the right side of the menu bar, select an object or enter the beginning characters of an object's name to narrow down the list of objects to select.

   After you select an object, the filter is applied and the list of rules is updated, showing only rules that contain the object in any of the following columns:

   - Sources
   - Destinations
   - Applied To
   - Services

4 To remove the filter, delete the object name from the text field.

Backing Up and Restoring NSX Manager or Global Manager

26

If an NSX Manager or a Global Manager appliance becomes inoperable, or if you want to restore your environment to a previous state, you can restore from a backup. NSX Managers are called Local Managers if they are managed with a Global Manager using NSX Federation.

You can restore an NSX configuration back to the captured state in any of the backups. While the appliance is inoperable, the data plane is not affected, but you cannot make configuration changes.

Note the following:

- For NSX backup authentication to connect to the SFTP backup server, you can use the user name and password method or have the NSX Manager use an SSH private key.
- You must restore the same version you have backed up from your NSX appliance, to the new NSX appliances.
- Use the same key size for backup and restore. If the key size is different at time of the backup and restore, the backup does not appear in the list of available backups. Support includes key size 256-bit, 384-bit, and 521-bit.
- NSX Manager or Global Manager restore can use the same IP or a different IP address.
  - If you use an NSX Manager or Global Manager IP address to restore, then you must use the same IP address as in the backup.
  - For managers with different IP addresses, you must configure an FQDN. If you use an NSX Manager or a Global Manager FQDN to restore, then you must use the same FQDN as in the backup. Use a lowercase FQDN only for backup and restore.
- If both the active and standby or only the active Global Managerfail, you need to:

Issue	Solution	Result
Both the active and standby fail.	1 Delete both active and standby Global Managers. Ensure there are no Global Manager appliances up in any other clusters. 2 Deploy a new Global Manager with the same IP address/FQDN as old active Global Manager. 3 Restore the active Global Manager from backup. If any Local Managers are present, they sync to the new active Global Manager. 4 Deploy a new Global Manager on another site and onboard it to the restore Global Manager.	■ The active Global Manager syncs with any Global Manager on the network. ■ The standby Global Manager syncs with the active Global Manager.
The active Global Manager fails and the standby Global Manager is up.	If standby Global Manager is in a good state, it automatically becomes the active Global Manager.	Standby Global Manager becomes active Global Manager. A new standby Global Manager must be manually added for backup. After the new active Global Manager is online, the Local Manager syncs up and ensures configuration replication.

This chapter includes the following topics:

- Configure Backups
- Start or Schedule Backups
- Remove Old Backups
- Listing Available Backups
- Restore a Backup

Configure Backups

Before backups can occur, you must configure a backup file server. After you configure a backup file server, you can start a backup at any time or schedule recurring backups. Administrators can choose to connect to the SFTP backup server using an SSH private key or a password-based authentication.

Prerequisites

- Verify that the SFTP server is running the supported OS and the supported SFTP software. The following table displays the supported and tested software for backup, although other software versions might work.

Currently supported OS	Specifically tested version	SFTP software version
CentOS	8.4	OpenSSH_8.0p1
	7.9 or 7.7	OpenSSH_7.4p1
RHEL	8.4	OpenSSH_8.0p1
	7.9 or 7.7	OpenSSH_7.4p1
Ubuntu	20.04	OpenSSH_8.2p1
	18.04	OpenSSH_7.6p1
Windows	Windows Server 2019 Standard	OpenSSH_for_Windows_8.1p1

- Verify that the SFTP server is ready for use and is running SSH and SFTP, using the following commands:
  - `$ ssh backup_user@sftp_server`
  - `$ sftp backup_user@sftp_server`
- Verify the required hashed ECDSA host key is present on the backup server. See Find the SSH Fingerprint of a Remote Server.
- Ensure that the directory path where you want to store your backups exists and that you have read/write permissions to that directory. You cannot use the root directory (/).
- If you are using the SSH private key option, ensure:
  - The SFTP backup server configuration includes an SSH public key linked to one of the users (in ~/.ssh/authorized_keys in a Linux server).
  - To complete the NSX Manager configuration, you have the corresponding private key.
  - You verify that you are storing the private key in a place other than the configuration.
- If you have multiple NSX deployments, you must use a different directory for storing the backup of each deployment.
- If your NSX Manager or your Global Manager appliance has the DNS server access set to "publish_fqdns": true, you must configure that setting on the new NSX Manager or new Global Manager appliance before restore. Follow instructions at "Configuring NSX Manager for Access by the DNS Server" in the *NSX Installation Guide*.

Procedure

1. From a browser, log in with admin privileges to the NSX Manager or Global Manager at https://<manager-ip-address>.
2. Select **System > Backup & Restore**.

3 Click **Edit** under the **SFTP Server** label to configure your SFTP server.

4 Enter the FQDN or IP address of the backup file server.

The protocol text box is already filled in. SFTP is the only supported protocol.

5 Change the default port if necessary. The default TCP port is 22.

6 In the **Directory Path** text box, enter the absolute directory path where the backups will be stored.

The directory must already exist and cannot be the root directory (`/`). Avoid using path drive letters or spaces in directory names; they are not supported. If the backup file server is a Windows machine, you must use the forward slash when you specify the destination directory. For example, if the backup directory on the Windows machine is `c:\SFTP_Root\backup`, specify `/SFTP_Root/backup` as the destination directory.

The path to the backup directory can contain only the following characters: alphanumerics (a-z , A-Z, 0-9), underscore (_) , plus and minus sign (+ -), tilde and percent sign (~ %), forward slash (/), and period (.).

The backup process generates a name for the backup file that can be quite long. On a Windows server, the length of the full path name of the backup file can exceed the limit set by Windows and cause backups to fail. To avoid this issue, see the KB article https://kb.vmware.com/s/article/76528.

7 Choose which authentication method you want to use to log into the backup file server.

a To enter a user name and password that authenticates to the backup file server, select **Password** and enter the required information.

b To use an SSH private key to send NSX backups to the SFTP server, select **SSH Private Key** and enter the required information.

If you edit the backup configuration, you do not need to re-enter the password or SSH Private Key.

8 You can leave the SSH Fingerprint blank and accept or reject the fingerprint provided by the server after you click **Save** in a later step. If necessary, you can retrieve the SSH fingerprint by using this API: `POST /api/v1/cluster/backups?action=retrieve_ssh_fingerprint`.

For more details, see Find the SSH Fingerprint of a Remote Server.

9 Verify that the required ECDSA host key is present on the backup server by running `#ssh-keyscan -t ecdsa <backup server IP/FQDN>`.

```
#ssh-keyscan -t ecdsa ftpserver.corp.local
   #ftpserver.corp.local:22 SSH-2.0-OpenSSH_7.6p1 Ubuntu-4ubuntu0.5
   ftpserver.corp.local ecdsa-sha2-nistp256
```

Starting in 4.1, NSX supports the ECDSA cipher, TLS_ECDHE_ECDSA_WITH_AES_256_SHA384, in support of EAL4 compliance. Support also includes RSA for SSH private key generation using key sizes 1024-bits, 2048-bits, and 4096-bits. 4096-bits is recommended. If the command output does not return a supported ECDSA key or if the key is empty, then the configured backup server does not have the supported cipher ECDSA and you must update the configuration and generate new keys on the backup server. Contact the OS vendor if you need guidance for that configuration.

10 Enter a passphrase.

Important You will need this passphrase to restore a backup. If you forget the passphrase, you cannot restore any backups.

11 Click **Save**.

Results

The Backup and Restore page refreshes with the newly configured SFTP server updated.

What to do next

After you successfully configure a backup file server, you can click **Start Backup** to manually start a backup immediately. Or, to schedule recurring automatic backups see Start or Schedule Backups. To see a list of available backups if you cannot access an NSX Manager or Global Manager appliance, see Listing Available Backups for details.

Start or Schedule Backups

After you configure your SFTP file server you can start a backup at any time or schedule recurring backups.

When you set up recurring backups, the system automatically backs up the inventory if there is an inventory change, such as the addition or removal of a Transport Node. This feature is not available for manual backups. You can also trigger backups for configuration changes. You can optionally select both options for recurring backups.

Inventory backups do not get collected for Global Manager.

Prerequisites

Complete the instructions at Configure Backups.

Procedure

1 To schedule a backup, click **Edit** under the **Schedule** label on the **System > Backup and Restore**.

   a Click the **Recurring Backup** toggle.

   b To set a weekly schedule, click **Weekly** and set the days and time of the backup.

   c To set up to a 24 hour interval, click **Interval** and set the interval between backups.

2 To trigger an unscheduled full configuration backup when the system detects any user, runtime, or non-configuration related changes, click the **Detect NSX configuration change** toggle.

   For Global Manager, this setting triggers a backup when the system detects any changes in the database, such as the addition or removal of a Local Manager, tier-0 gateway, or DFW policy.

   You can specify a time interval for detecting database configuration changes. The valid range is 5 minutes to 1,440 minutes (24 hours). This option can potentially generate a large number of backups. Use it with caution.

Results

If you selected **Start Backup**, you see a progress bar of the in-progress backup.

When the manual or scheduled backup completes, the backup gets listed in the Backup History section of the page. The **Last Backup Status** label indicates whether the backup was successful and lists the timestamp, node, and cluster details of the appliance backed up. If the backup fails, you can see an error message.

What to do next

To restore a backup follow instructions at Restore a Backup.

Remove Old Backups

Backups can accumulate on the backup file server and consume a large amount of storage. You can run a script that comes with NSX to automatically delete old backups.

You can find the Python script `nsx_backup_cleaner.py` in the directory `/var/vmware/nsx/file-store` on NSX Manager. To access this file, you must log in as root. Typically, you schedule a job on the backup file server to run this script periodically to clean up old backups. The script works only for subfolders named cluster-node-backups and inventory-summary. The script fails if any other subfolders are present in the backup directory other than cluster-node-backups and inventory-summary.

The following usage information describes how to run the script:

```
nsx_backup_cleaner.py -d backup_dir [-k 1] [-l 5] [-h]
Or
nsx_backup_cleaner.py --dir backup_dir [--retention-period 1] [--min-count 5] [--help]

Required parameters:
    -d/--dir: Backup root directory
    -k/--retention-period: Number of days need to retain a backup file

Optional parameters:
    -l/--min-count: Minimum number of backup files to be kept, default value is 100
    -h/--help: Display help message
```

The age of a backup is calculated as the difference between the backup's timestamp and the time the script is run. If this value is larger than the retention period, the backup is deleted if there are more backups on the disk than the minimum number of backups.

For more information about setting up the script to run periodically on a Linux or Windows server, see the comments at the beginning of the script.

Listing Available Backups

The backup file server stores backups from all the NSX Manager or Global Manager nodes. To get the list of backups so that you can find the one you want to restore, you must run the `get_backup_timestamps.sh` script.

Note Use the same key size for backup and restore. If the key size is different at time of backup and restore, the backup does not appear in the list. Starting in NSX 3.2.1, support includes key size 256-bit, 384-bit, and 521-bit. In 3.2.0, support includes only 256-bit key size.

The script can be found on each NSX Manager or Global Manager appliance at `/var/vmware/nsx/file-store/get_backup_timestamps.sh`. You can run this script on any Linux machine or on the NSX appliance. As a best practice, copy this script after installing NSX to a machine that is not an NSX Manager or Global Manager so that you can run this script even if all the NSX Manager or Global Manager nodes become inaccessible. If you need to restore a backup but have no access to this script, you can install a new NSX Manager or Global Manager node and run the script there.

You can copy the script to another machine or to the backup file server by logging in to the NSX Manager or Global Manager as admin and running a CLI command. For example:

```
nsxmgr-1> copy file get_backup_timestamps.sh url scp://admin@server1/tmp/
admin@server's password:
nsxmgr-1>
```

The script is interactive and prompts you for the information that you specified when you configured the backup file server. You can specify the number of backups to display. Each backup is listed with a timestamp, the NSX Manager or Global Manager node's IP address or FQDN if the NSX Manager or Global Manager node is set up to publish its FQDN, and the node ID. For example,

```
admin@host1:/home/admin# ./get_backup_timestamps.sh
Enter file server ip:
10.10.10.20
Enter port:
22
Enter directory path:
/home/nsx/backups
Enter number of latest backup or press Enter to list all backups:

root@10.10.10.20's password:
Latest backups:
```

```
[Backup timestamp; IP address/FQDN; Node id]
2019-01-22;09:16:43 nsxmgr.example.com 41893642-597b-915f-5117-7da576df4ff2
2019-01-22;09:14:42 nsxmgr.example.com 41893642-597b-915f-5117-7da576df4ff2
2019-01-22;09:13:30 nsxmgr.example.com 41893642-597b-915f-5117-7da576df4ff2
2019-01-22;09:01:52 10.10.10.77 35163642-6623-8f6d-7af0-52e03f16faed
2019-01-22;09:00:33 10.10.10.77 35163642-6623-8f6d-7af0-52e03f16faed
```

Restore a Backup

Restoring a backup restores the state of the network at the time of the backup. In addition, the configurations maintained by NSX Manager or Global Manager appliances are also restored. For NSX Manager, any changes, such as adding or deleting nodes, that were made to the fabric since the backup was taken, are reconciled. NSX Managers are called Local Managers (LM) when they are federated with a Global Manager (GM).

Note DNS entries (name servers and search domains) are not retained when you restore from a backup. To redeploy in a VMware Cloud Foundation (VCF) deployment using an OVF file, you must use FQDNs for the NSX Manager VM names.

You must restore the backup to a new NSX Manager or Global Manager appliance. Follow the instructions for your specific case.

- If you had a cluster of the NSX Manager appliance when the backup was taken, the restore process restores one node first and then prompts you to add the other nodes. You can add the other nodes during the restore process or after the first node is restored. See the following detailed steps.
- If you had a cluster of Global Manager appliances, you can only restore one node using the restore process. You must create the cluster after the restore of the first node completes. For instructions on restoring a lost active Global Manager, a lost standby Global Manager, or a lost Local Manager, see Backup and Restore in NSX Federation.

Important If any nodes in the appliance cluster are still available, you must power them off before you start the restore.

Prerequisites

- Verify that you have the login credentials (username and password or the SSH private key matching the public key stored on the backup file server) for the backup file server.
- Verify that you have the SSH fingerprint of the backup file server. Support includes key size 256-bit, 384-bit, and 521-bit. Ensure whatever key size is used at time of backup is used at time of restore.
- Verify that you have the passphrase of the backup file.
- Identify which backup you want to restore by following the procedure in Listing Available Backups. Take note of the IP or FQDN of the NSX appliance that took the backup.

- Ensure the network setup where you are performing the restore has the same set of network connectivity as the system on which you performed the backup. For example, the same VIPs, DNS, NTP communication,and so on. If network connectivity is not same, fix the inconsistencies before adding a second or third node to the restored system.
- Perform a federated restore when both the active and standby Global Managers are down. If this is not the case, see Backup and Restore in NSX Federation.

Procedure

1 If any nodes in the appliance cluster are still available, you must power them off before you start the restore.

2 Install one new appliance node on which to restore the backup.

   - If the backup listing for the backup you are restoring contains an IP address, you must deploy the new NSX Manager or Global Manager node with the same IP address. Do not configure the node to publish its FQDN.
   - If the backup listing for the backup you are restoring contains an FQDN, you must configure the new appliance node with this FQDN and publish the FQDN. Only lowercase FQDN is supported for backup and restore.

     Note Until the FQDN is configured and published, the Restore button for the backup is disabled in the newly deployed NSX Manager or Global Manager UI.

     Use this API to publish the NSX Manager or Global Manager FQDN.

     Example request:

     ```
     PUT https://<nsx-mgr OR global-mgr>/api/v1/configs/management

     {
       "publish_fqdns": true,
       "_revision": 0
     }
     ```

     See the *NSX API Guide* for API details.

     In addition, if the new manager node has a different IP address than the original one, you must update the DNS server's forward and reverse lookup entries for the manager node with the new IP address.

   After the new manager node is running and online, you can proceed with the restore.

3 From a browser, log in with admin privileges to the NSX Manager or Global Manager at https://<manager-ip-address>.

4 Select **System > Backup & Restore**.

5 To configure the backup file server, click **Edit**.

   Do not configure automatic backup if you are going to perform a restore.

6 Enter the IP address or FQDN.

7 Change the port number, if necessary.

The default is 22.

8 In the **Directory Path** text box, enter the absolute directory path where the backups are stored.

The path to the backup directory can contain only the following characters: alphanumerics (a-z , A-Z, 0-9), underscore (_) , plus and minus sign (+ -), tilde and percent sign (~ %), forward slash (/), and period (.).

Avoid using path drive letters or spaces in directory names; they are not supported. If the backup file server is a Windows machine, you must use the forward slash when you specify the destination directory. For example, if the backup directory on the Windows machine is c:\SFTP_Root\backup, specify /SFTP_Root/backup as the destination directory.

9 To log in to the server, enter the user name and password or the SSH private key, depending on your backup authentication scheme.

10 You can leave the SSH Fingerprint blank and accept or reject the fingerprint provided by the server after you click **Save** in a later step. If necessary, you can retrieve the SSH fingerprint by using this API: `POST /api/v1/cluster/backups?action=retrieve_ssh_fingerprint`.

11 Enter the passphrase that was used to encrypt the backup data.

12 Click **Save**.

13 Select a backup.

14 Click **Restore**.

15 The restore process prompts you to take action, if necessary, as it progresses.

Note If you are restoring a Global Manager appliance, the following steps do not appear. After restoring the first Global Manager node, you must manually join the other nodes to form the cluster. If you are restoring a multi-site network, see the "Limitations" section of the NSX Multisite topic.

a Confirm CM/VC Connectivity: If you want to restore existing compute managers, ensure that they are registered with the new NSX Manager node and available during the restore process.

b If you have deleted or added fabric nodes or transport nodes, you are prompted to take certain actions, for example, log in to a node and run a script. If you have created a logical switch or segment since the backup, the logical switch or segment will not appear after the restore.

c If the backup has information about a manager cluster, you are prompted to add other nodes. If you decide not to add nodes, you can still proceed with the restore and manually add other nodes to form the cluster after the restore of this node completes.

d If there are fabric nodes that did not discover the new manager node, you are provided a list of them.

A progress bar displays the status of the restore operation noting the step the restore process is on. During the restore process, services on the manager appliance get restarted and the control plane becomes unavailable until restore completes.

After the restore operation is finished, the **Restore Complete** screen shows the result of the restore, the timestamp of the backup file, and the start and end time of the restore operation. Any segments created after the backup was taken are not restored.

If the restore fails, the screen displays the step where the failure occurred, for example, `Current Step: Restoring Cluster (DB)` or `Current Step: Restoring Node`. If either cluster restore or node restore fails, the error might be transient. In that case, there is no need to click **Retry**. You can restart or reboot the manager and the restore continues.

You can also determine if there was a cluster or node restore failure by selecting the log files. Run `get log-file syslog` to view the system log file and search for the strings `Cluster restore failed` and `Node restore failed`.

To restart the manager, run the `restart service manager` command.

To reboot the manager, run the `reboot` command.

Note If you added a compute manager after the backup, and you try to add the compute manager again after the restore, you get an error message indicating that registration failed. Click the **Resolve** button to resolve the error and successfully add the compute manager. For more information, see Add a Compute Manager, step 4. If you want to remove information about NSX that is stored in a VMware vCenter, follow the steps in Remove NSX Extension from VMware vCenter

If the VMware vCenter was registered with custom ports in the backup, you must manually open all the custom ports on the restored manager appliances.

16 If you have only one node deployed, after the restored manager node is up and functional, you can deploy additional nodes to form a cluster.

See the *NSX Installation Guide* for instructions.

17 If you had other manager cluster VMs that you powered down in Step 1, delete them after the new manager cluster is deployed.

Operations and Management

27

You may need to change the configuration of the appliances you've installed, for example, adding licenses, certificates, and changing passwords. There are also routine maintenance tasks that you should perform, including running backups. Additionally, there are tools to help you find information about the appliances that are part of the NSX infrastructure and the logical networks created by NSX, including remote system logging, traceflow, and port connections.

This chapter includes the following topics:

- View the Usage and Capacity of Categories of Objects
- Configuring the Login Banner and UI
- Configure a Node Profile
- Checking the Realized State of a Configuration Change
- View Network Topology
- Search for Objects
- Filter by Object Attributes
- Add a Compute Manager
- Configuring Active Directory and Event Log Scraping
- Enable Windows Security Log Access for the Event Log Reader
- Add an LDAP Server
- Synchronize Active Directory
- Remove NSX Extension from VMware vCenter
- Managing the NSX Manager Cluster
- Replacing an NSX Edge Transport Node in an NSX Edge Cluster
- Managing Resource Reservations for an Edge VM Appliance
- Replacing NSX Edge Hardware or Redeploying NSX Edge Nodes VM
- Configure Advanced NSX Edge Parameters to Improve Performance
- Edit NSX Edge VM Restart Policy
- Adding and Removing an ESXi Host Transport Node to and from vCenter Servers

- Changing the Distributed Router Interfaces' MAC Address
- Configuring Appliances
- Configuring NTP on Appliances and Transport Nodes
- Add a License Key and Generate a License Usage Report
- Standards Compliance
- Collect Support Bundles
- Understanding Support Bundle File Paths
- Debugging NSX at Runtime
- Log Messages and Error Codes
- Customer Experience Improvement Program
- Find the SSH Fingerprint of a Remote Server
- Configuring an External Load Balancer
- Configure Proxy Settings
- Promote Manager Objects to Policy Objects
- Back up and restore NSX configured in VMware vCenter
- Configure an External Entropy Source

View the Usage and Capacity of Categories of Objects

You can view the usage and capacity of different categories of objects. By default, an alarm is generated when the current inventory a category of objects reaches a certain level.

You can set thresholds for each object category. If a threshold is reached, an alarm will be generated. You can see alarms by navigating to **Home > Alarms**.

Note If both manager and policy objects exist, only the capacity information for manager objects will be available. Also, if you configure the user interface settings and make the **Policy/Manager** toggle visible, the capacity information for policy objects will not be available even if there are no manager objects. Capacity information for policy objects will only be shown if there are no manager objects and the **Policy/Manager** toggle is not visible.

To see the usage and capacity of different categories of objects, click one of the following tabs:

- **Networking > Network Overview > Capacity**
- **Security > Security Overview > Capacity**
- **Inventory > Inventory Overview > Capacity**
- **System > System Overview > Capacity**

You can also navigate to **Plan & Troubleshoot > Consolidated Capacity** to see all the object categories on one page.

On each capacity page, for each category of objects, the following information is displayed:

- Current Inventory - The number of objects that have been successfully created or configured. A color-coded bar is displayed to indicate the usage percentage. If usage is below the minimum capacity threshold, the color is green. If usage is at or above the minimum capacity threshold but below the maximum capacity threshold, the color is orange. If usage is at or above the maximum capacity threshold, the color is red.
- Maximum Capacity.
- Minimum Capacity Threshold - This is the usage level at which the usage bar mentioned above will show an orange color. You can change this value. The default is 70%.
- Maximum Capacity Threshold - This is the usage level at which the usage bar mentioned above will show a red color. You can change this value. The default is 100%.

If you change the minimum or maximum capacity threshold, you can click **Revert** to go back to the last saved value. You can click **Reset Values** to restore the default values for all the object categories.

Object Categories in Policy Mode

The networking capacity page shows the following object categories:

- Segments
- System-wide DHCP Ranges
- Tier-0 Gateways
- Tier-1 Gateways
- DHCP Server Instances
- System-wide NAT Rules
- Prefix Lists
- Tier-1 Gateways with NAT Enabled
- Segment Ports

The security capacity page shows the following object categories:

- Introspection Rules N-S Tier-1
- Service Chains
- Active Directory Groups (Identity Firewall)
- Saved Firewall Rules Configuration
- Introspection Policies E-W
- Introspection Policies N-S Tier-0

- System-wide Firewall Rules
- System-wide Endpoint Protection Enabled Virtual Machines
- Introspection Rules N-S Tier-0
- Introspection Rules E-W
- Distributed Firewall Sections
- Introspection Policies N-S Tier-1
- System-wide Firewall Sections
- Active Directory Domains (Identity Firewall)
- System-wide Endpoint Protection Enabled Hosts
- Distributed Firewall Rules
- Service Paths

The inventory capacity page shows the following object categories:

- Groups
- vSphere Clusters
- Services
- Hypervisor Hosts
- Groups based on IP Addresses

The system capacity page shows the following object categories:

- Compute Managers
- System-wide Edge Nodes
- Edge clusters

Object Categories in Manager Mode

The networking capacity page shows the following object categories:

- Tier-0 logical routers
- Tier-1 logical routers
- Prefix lists
- System-wide NAT rules
- DHCP server instances
- System-wide DHCP ranges and pools
- Tier-1 logical routers with NAT enabled
- Logical switches

- System-wide logical switch ports

The security capacity page shows the following object categories:

- System-wide endpoint protection-enabled hosts
- System-wide endpoint protection-enabled virtual machines
- Active Directory groups
- Active Directory domains
- Distributed firewall rules
- System-wide firewall rules
- System-wide firewall sections
- Distributed firewall sections

The inventory capacity page shows the following object categories:

- Groups
- IP sets
- Groups based on IP sets
- vSphere clusters
- Hypervisor hosts

The system capacity page shows the following object categories:

- Edge clusters
- System-wide edge nodes

Configuring the Login Banner and UI

You can configure a custom banner for the login window and customize your user interface for NSX Managers, Global Managers, and Local Managers.

Configure the Login Window with a User Agreement Banner

You can configure a custom login message that requires a user agreement as part of the login process for NSX.

If activated (or enabled), this agreement displays at every login and is changed only by users with Enterprise admin roles. For Federation, each Global and Local Manager requires their own configuration. If you have DoD or other security compliance requirements, this feature provides such options.

Procedure

1 From your browser, log in with admin privileges to an NSX Manager or Global Manager at `https://<nsx-manager-ip-address>`.

2 Navigate to **System > General System Settings**.

3 To add or remove banner settings, click the **User Interface** tab and click **Edit** for **Login Consent Settings**.

4 To activate or deactivate this option, toggle the Login Consent **On** or **Off**.

5 To require explicit user consent before the user logs in, select **Yes**.

6 To create a consent message that displays on the login page, enter the consent message title.

   This message starts with the phrase **I agree to …** and can be up to 255 characters.

7 If required, add a custom consent message description.

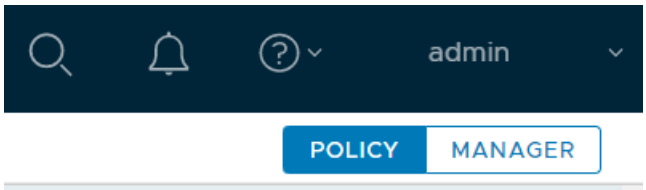
   The description limit is 64,000 characters. Users must select the title to view this description. HTML links are not supported.

Configure the User Interface Settings

You can configure how your users view the NSX user interface. These settings are valid for NSX Manager, and in NSX Federation for Global Managers and Local Managers.

Prior to NSX 3.2, users could access two possible modes in the user interface: Policy and Manager. You can control which mode is default, and whether users can switch between them using the user interface mode buttons. The Policy mode is the default. New users of release 4.0 will not see the **Manager** button.

If present, you can use the **Policy** and **Manager** buttons to switch between the Policy and Manager modes. Switching modes controls which menus items are available to you.



- By default, if your environment contains only objects created through Policy mode, your user interface is in Policy mode and you do not see the **Policy** and **Manager** buttons.
- By default, if your environment contains any objects created through Manager mode, you see the **Policy** and **Manager** buttons in the top-right corner.

You can use the User Interface settings to modify these defaults.

See Chapter 1 NSX Manager for more information about the modes.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Navigate to **System > General System Settings**.

3 Select **User Interface** and click **Edit** in the User Interface Mode Toggle pane.

4 Modify the user interface settings: **Toggle Visibility** and **Default Mode**.

Toggle Visibility	Description
Visible to All Users	If Manager mode objects are present, the mode buttons are visible to all users.
Visible to Users with the Enterprise Admin Role	If Manager mode objects are present, the mode buttons are visible to users with the **Enterprise Admin** role.
Hidden from All Users	Even if Manager mode objects are present, the mode buttons are hidden from all users. This displays Policy mode UI only, even if Manager mode objects are present.
Default Mode	Can be set to Policy or Manager, if available.

Configure a Node Profile

You can configure settings such as time zone, NTP servers, SNMP, and syslog servers to apply to all NSX Manager and Edge nodes.

In this release, only one node profile is supported. This profile represents a collection of time zone, NTP servers, SNMP configuration and syslog servers. By default, the node profile is applied to all nodes, unless the node is configured to not accept such configuration from the NSX Manager. To prevent a node from accepting the node profile, use the CLI command `set node central-config disabled` on that node.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **System > Fabric > Profiles**.

3 Click the **Node Profiles** tab.

4 Click **All NSX Nodes** in the **Name** column.

5 Click **Edit** to configure the time zone and NTP servers.

6 In the **Syslog Servers** section, click **Add** to add a Syslog server.

   a Enter the FQDN or IP address of the Syslog server.

   b Specify a port number.

c Select a protocol.

The available protocols are **TCP**, **UDP**, and **LI (Aria Operations for Logs)**.

d Select a log level.

The available levels are **Emergency**, **Alert**, **Critical**, **Error**, **Warning**, **Notice**, **Information**, and **Debug**.

7 In the **SNMP Polling** section, under **v2c**, click **Add** to add an SNMPv2c community.

a Enter a name for the community.

b Enter a **Community String** value.

This value is used for authentication.

8 In the **SNMP Polling** section, under **v3**, click **Add** to add an SNMPv3 user.

a Enter a user name.

b Enter an authentication password.

You can click the icon on the right to show or hide the password.

c Enter a private password.

You can click the icon on the right to show or hide the password.

9 In the **SNMP Traps** section, under **v2c**, click **Add** to add an SNMPv2c trap configuration.

a Enter a FQDN or IP address.

b Specify a port number.

c Enter a name for the community.

d Enter a **Community String** value.

This value is used for authentication.

10 In the **SNMP Traps** section, under **v3**, click **Add** to add an SNMPv3 trap configuration.

a Enter a FQDN or IP address.

b Specify a port number.

c Enter a user name.

What to do next

Verify that the profile configurations are applied to the NSX Manager and NSX Edge nodes. Log in to the NSX Manager and NSX Edge nodes with **admin** privileges, and run the following commands:

- `get clock`
- `get ntp-server`
- `get logging-servers`

- `get snmp v2-targets`
- `get snmp v3-targets`
- `get snmp v2-configured`
- `get snmp v3-configured`
- `get snmp v3-engine-id`
- `get snmp v3-protocols`
- `get snmp v3-users`

For more information about these commands including examples, see the *NSX Command-Line Interface Reference*.

Error situations

If the node profile configurations are not applied successfully, then there are two possibilities:

- The central configuration was not synchronized with the remote node due to connectivity issues between NSX Manager and the remote node. In this case, you cannot do anything from the central configuration side.
- The central configuration was synchronized with the remote node, but the command to apply the central configuration failed to run. In this case, you can check syslog on the remote node.

  In the logs, search for the `subcomp="central_node_config_update"` string to look for any errors.

  For example, the syslog exporter configuration might fail if the host name specified cannot be resolved to IP addresses, or if a second Aria Operations for Logs server is being configured.

  The following example logs show the error messages:

  Log example 1:

```
2020-05-18T22:56:06.485Z vmw-svc.nsxmanager-sb-36265022-1-rhel NSX 24904 - [nsx@6876
comp="nsx-manager" subcomp="central_node_config_update" username="root" level="INFO"]
No change in timezone 2020-05-18T22:56:07.184Z vmw-svc.nsxmanager-sb-36265022-1-
rhel NSX 24904 - [nsx@6876 comp="nsx-manager" subcomp="central_node_config_update"
username="root" level="INFO"] No change in NTP configuration 2020-05-18T22:56:07.210Z
vmw-svc.nsxmanager-sb-36265022-1-rhel NSX 24904 - [nsx@6876 comp="nsx-manager"
subcomp="central_node_config_update" username="root" level="INFO"] Updating Syslog
configuration 2020-05-18T22:56:08.826Z vmw-svc.nsxmanager-sb-36265022-1-rhel NSX 24904
- [nsx@6876 comp="nsx-manager" subcomp="central_node_config_update" username="root"
level="WARNING"] Failed to add syslog exporter {"port": 514, "exporter_name":
"264aa005-dfb0-4942-a1c4-f749bfc1a2c4", "protocol": "TCP", "level": "ERR", "server":
"vikas.2020.com"}, response: {#012  "error_code": 36569,#012  "error_message": "Error
modifying firewall rule due to invalid hostname.",#012  "module_name": "node-
services"#012}, status: 400, err: 400 Client Error: Bad Request for url: http://
localhost:7441/api/v1/node/services/syslog/exporters
```

Log example 2:

```
2020-05-18T22:56:08.839Z vmw-svc.nsxmanager-sb-36265022-1-rhel NSX 24904 -
[nsx@6876 comp="nsx-manager" subcomp="central_node_config_update" username="root"
level="WARNING"] Failed to add syslog exporter {"port": 514, "exporter_name":
"f4e088d4-4b45-42fe-ba1d-7f98838c7f61", "protocol": "LI", "level": "INFO", "server":
"loginsight.vmware.com"}, response: {#012  "error_code": 36400,#012  "error_message":
"Maximum number of loginsight servers exceeded",#012  "module_name": "node-
services"#012}, status: 400, err: 400 Client Error: Bad Request for url: http://
localhost:7441/api/v1/node/services/syslog/exporters
```

Log example 3:

```
2020-05-18T22:56:10.639Z vmw-svc.nsxmanager-sb-36265022-1-rhel NSX 24904 -
[nsx@6876 comp="nsx-manager" subcomp="central_node_config_update" username="root"
level="WARNING"] Failed to add syslog exporter {"port": 514, "exporter_name":
"d0dc1797-b5dc-42ba-b07d-fe107dd70111", "protocol": "UDP", "level": "INFO", "server":
"logging.vmware.com"}, response: {#012  "error_code": 36569,#012  "error_message":
"Error modifying firewall rule due to invalid hostname.",#012  "module_name": "node-
services"#012}, status: 400, err: 400 Client Error: Bad Request for url: http://
localhost:7441/api/v1/node/services/syslog/exporters
```

Checking the Realized State of a Configuration Change

When you make a configuration change, NSX Manager typically sends a request to another component to implement the change. For some layer 3 entities, if you make the configuration change using the API, you can track the status of the request to see if the change is successfully implemented.

The configuration change that you initiate is called the desired state. The result of implementing the change is called the realized state. If NSX Manager implements the change successfully, the realized state will be the same as the desired state. If there is an error, the realized state will not be the same as the desired state.

For some layer 3 entities, when you call an API to make a configuration change, the response will include the parameter `request_id`. You can use the parameters `request_id` and the `entity_id` to make an API call to find out the status of the request.

This feature supports the following entities and APIs:

```
EdgeCluster
    POST /edge-clusters
    PUT /edge-clusters/<edge-cluster-id>
    DELETE /edge-clusters/<edge-cluster-id>
    POST /edge-clusters/<edge-cluster-id>?action=replace_transport_node

LogicalRouter
    POST /logical-routers
    PUT /logical-routers/<logical-router-id>
    DELETE /logical-routers/<logical-router-id>
    POST /logical-routers/<logical-router-id>?action=reprocess
    POST /logical-routers/<logical-router-id>?action=reallocate
```

```
LogicalRouterPort
    POST /logical-router-ports
    PUT /logical-router-ports/<logical-router-port-id>
    DELETE /logical-router-ports/<logical-router-port-id>

StaticRoute
    POST /logical-routers/<logical-router-id>/routing/static-routes
    PUT /logical-routers/<logical-router-id>/routing/static-routes/<static-route-id>
    DELETE /logical-routers/<logical-router-id>/routing/static-routes/<static-route-id>

BGPConfig
    PUT /logical-routers/<logical-router-id>/routing/bgp

BgpNeighbor
    POST /logical-routers/<logical-router-id>/routing/bgp/neighbors
    PUT /logical-routers/<logical-router-id>/routing/bgp/neighbors/<bgp-neighbor-id>
    DELETE /logical-routers/<logical-router-id>/routing/bgp/neighbors/<bgp-neighbor-id>
    POST /logical-routers/<logical-router-id>/routing/bgp/neighbors/<bgp-neighbor-id>

BGPCommunityList
    POST /logical-routers/<logical-router-id>/routing/bgp/community-lists
    PUT /logical-routers/<logical-router-id>/routing/bgp/community-lists/<community-list-id>
    DELETE /logical-routers/<logical-router-id>/routing/bgp/community-lists/<community-list-
id>

AdvertisementConfig
    PUT /logical-routers/<logical-router-id>/routing/advertisement

AdvertiseRouteList
    PUT /logical-routers/<logical-router-id>/routing/advertisement/rules

NatRule
    POST /logical-routers/<logical-router-id>/nat/rules
    PUT /logical-routers/<logical-router-id>/nat/rules/<rule-id>
    DELETE /logical-routers/<logical-router-id>/nat/rules/<rule-id>

DhcpRelayService
    POST /dhcp/relays
    PUT /dhcp/relays/<relay-id>
    DELETE /dhcp/relays/<relay-id>

DhcpRelayProfile
    POST /dhcp/relay-profiles
    PUT /dhcp/relay-profiles/<relay-profile-id>
    DELETE /dhcp/relay-profiles/<relay-profile-id>

StaticHopBfdPeer
    POST /logical-routers/<logical-router-id>/routing/static-routes/bfd-peers
    PUT /logical-routers/<logical-router-id>/routing/static-routes/bfd-peers/<bfd-peers-id>
    DELETE /logical-routers/<logical-router-id>/routing/static-routes/bfd-peers/<bfd-peers-id>

IPPrefixList
    POST /logical-routers/<logical-router-id>/routing/ip-prefix-lists
    PUT /logical-routers/<logical-router-id>/routing/ip-prefix-lists/<ip-prefix-list-id>
```

```
    DELETE /logical-routers/<logical-router-id>/routing/ip-prefix-lists/<ip-prefix-list-id>

RouteMap
    POST /logical-routers/<logical-router-id>/routing/route-maps
    PUT /logical-routers/<logical-router-id>/routing/route-maps/<route-map-id>
    DELETE /logical-routers/<logical-router-id>/routing/route-maps/<route-map-id>

RedistributionConfig
    PUT /logical-routers/<logical-router-id>/routing/redistribution
RedistributionRuleList
    PUT /logical-routers/<logical-router-id>/routing/redistribution/rules

BfdConfig
    PUT /logical-routers/<logical-router-id>/routing/bfd-config

MplsConfig
    PUT /logical-routers/<logical-router-id>/routing/mpls

RoutingGlobalConfig
    PUT /logical-routers/<logical-router-id>/routing

IPSecVPNIKEProfile
    POST /vpn/ipsec/ike-profiles
    PUT /vpn/ipsec/ike-profiles/<ike-profile-id>
    DELETE /vpn/ipsec/ike-profiles/<ike-profile-id>

IPSecVPNDPDProfile
    POST /vpn/ipsec/dpd-profiles
    PUT /vpn/ipsec/dpd-profiles/<dpd-profile-id>
    DELETE /vpn/ipsec/dpd-profiles/<dpd-profile-id>

IPSecVPNTunnelProfile
    POST /vpn/ipsec/tunnel-profiles
    PUT /vpn/ipsec/tunnel-profiles/<tunnel-profile-id>
    DELETE /vpn/ipsec/tunnel-profiles/<tunnel-profile-id>

IPSecVPNLocalEndpoint
    POST /vpn/ipsec/local-endpoints
    PUT /vpn/ipsec/local-endpoints/<local-endpoint-id>
    DELETE /vpn/ipsec/local-endpoints/<local-endpoint-id>

IPSecVPNPeerEndpoint
    POST /vpn/ipsec/peer-endpoints
    PUT /vpn/ipsec/peer-endpoints/<peer-endpoint-id>
    DELETE /vpn/ipsec/peer-endpoints/<peer-endpoint-id>

IPSecVPNService
    POST /vpn/ipsec/services
    PUT /vpn/ipsec/services/<service-id>
    DELETE /vpn/ipsec/services/<service-id>

IPSecVPNSession
    POST /vpn/ipsec/sessions
    PUT /vpn/ipsec/sessions/<session-id>
    DELETE /vpn/ipsec/sessions/<session-id>
```

```
DhcpServer
    POST /dhcp/servers
    PUT /dhcp/servers/<server-id>
    DELETE /dhcp/servers/<server-id>

DhcpStaticBinding
    POST /dhcp/servers/static-bindings
    PUT /dhcp/servers/<server-id>/static-bindings/<binding-id>
    DELETE /dhcp/servers/<server-id>/static-bindings/<binding-id>

DhcpIpPool
    POST /dhcp/servers/ip-pools
    PUT /dhcp/servers/<server-id>/ip-pools/<pool-id>
    DELETE /dhcp/servers/<server-id>/ip-pools/<pool-id>

DnsForwarder
    POST /dns/forwarders
    PUT /dns/forwarders/<forwarder-id>
    DELETE /dns/forwarders/<forwarder-id>
```

You can call the following APIs to get the realized states:

```
EdgeCluster
Request - GET /edge-clusters/<edge-cluster-id>/state?request_id=<request-id>
Response - An instance of EdgeClusterStateDto which will inherit ConfigurationState. If the
edge cluster is deleted then the state will be unknown and it will return the common entity
not found error.

LogicalRouter / All L3 Entites - All L3 entities can use this API to get realization state
Request - GET /logical-routers/<logical-router-id>/state?request_id=<request-id>
Response - An instance of LogicalRouterStateDto which will inherit ConfigurationState. Delete
operation of any entity other than logical router can be covered by getting the state of
logical router but if the logical router itself is deleted then the state will be unknown and
it will return the common entity not found error.

LogicalServiceRouterCluster - All L3 entities which are the part of services can use this API
to get the realization state
Request - GET /logical-routers/<logical-router-id>/service-cluster/state?request_id=<request-
id>
Response - An instance of LogicalServiceRouterClusterState which will inherit
ConfigurationState.

LogicalRouterPort / DhcpRelayService / DhcpRelayProfile
Request - GET /logical-router-ports/<logical-router-port-id>/state?request_id=<request-id>
Response - An instance of LogicalRouterPortStateDto which will inherit ConfigurationState.

IPSecVPNIKEProfile / IPSecVPNDPDProfile / IPSecVPNTunnelProfile / IPSecVPNLocalEndpoint /
IPSecVPNPeerEndpoint / IPSecVPNService / IPSecVPNSession
Request - GET /vpn/ipsec/sessions/<session-id>/state?request_id=<request-id>
Response - An instance of IPSecVPNSessionStateDto which will inherit ConfigurationState. If
the session is deleted then the state will be unknown and it will return the common entity
not found error. When IPSecVPNService is disabled, IKE itself is down and it does not
respond. It will return unknown state in such a case.
```

```
DhcpServer
Request - GET /dhcp/servers/<server-id>/state?request_id=<request-id>
Response - An instance of ConfigurationState.

DhcpStaticBinding
Request - GET /dhcp/servers/<server-id>/static-bindings/<binding-id>/state?
request_id=<request-id>
Response - An instance of ConfigurationState.

DhcpIpPool
Request - GET /dhcp/servers/<server-id>/ip-pools/<pool-id>/state?request_id=<request-id>
Response - An instance of ConfigurationState.

DnsForwarder
Request - GET /dns/forwarders/<forwarder-id>/state?request_id=<request-id>
Response - An instance of ConfigurationState.
```

For more information about the APIs, see the *NSX API Guide*.

View Network Topology

View the network topology of your NSX environment for an overview of the entities, services, and the underlying fabric in your network. The graphical representation of the network topology is helpful when you are verifying your network configuration or troubleshooting errors.

Table 27-1. Entities and Services

Entity	Service
Network Services	Load Balancer, NAT, and Distributed Firewall
Tier-0	L2 VPN service, IPSec VPN service, NAT rules, DNS forwarder, Gateway firewall rules, DHCP
Tier-1	Load Balancer, L2 VPN service, IPSec VPN service, NAT rules, DNS forwarder, Gateway firewall rules, DHCP
Segment	DHCP, Metadata proxy
Edge Transport Node	Switches, uplink interfaces, host configuration, uplinks, PortGroup connectivity
VirtualMachine	Underlying host cluster and host transport nodes
Pod	None
PhysicalServer	None
VRF	NAT rules, Gateway firewall rules, DHCP

Procedure

1. From your browser, log in with admin privileges to an NSX Manager at https://<nsx-manager-ip-address>.
2. Select **Networking > Network Topology**.

3 Navigate through the network topology to see more information:

- Network connectivity - Zoom in to view more details about the logical entities. You can point to an entity to view its logical path in the network and view the services configured on it.
- Service configuration - Click a service to view more configuration details.

  For the IPSec VPN service, NSX displays a visual representation of configuration information like local and remote endpoints and also indicates whether the sessions are policy-based or rule-based. You can then click a session to view session-specific details like tunnel status, among others.
- Fabric view - Double-click an object to see the fabric view for that object and the parent object in its logical path.
  - For VMs and Containers, fabric view displays the host cluster name, host transport node details, and the configuration view for the host transport node.
  - For a physical server, fabric view displays details of the host transport node and the configuration view for the node.
  - For a tier-0 and a tier-1 gateway, fabric view displays the distributed router(DR), Edge cluster, Edge nodes where the service router(SR) is realized with HA status of the SR. You can also view configuration details of the edge node where the SR is realized for the gateway. For a tier-0 gateway, the uplink interfaces are also displayed.

    Click an Edge cluster to see a list of Edge transport nodes that are members of the cluster. You can click a node and expand **Tunnels** for additional information on the Edge tunnels BFD status using filters. Use the **Configuration** option to download routing or forwarding tables.

    Zoom in for the **View Edge Node Configuration** option to see a visual representation of the Edge node configuration.
- Click **Export** on the tool bar to save the topology to a PDF file.
- Apply filters to focus on specific objects. See Filter by Object Attributes for more details about filters.

Search for Objects

You can search for objects using various criteria throughout the NSX inventory.

The search results are sorted by relevance and you can filter these results based on your search query.

Note If you have special characters in your search query that also function as operators, then you must add a leading backslash. The characters that function as operators are: +, -, =, &&, ||, <, >, !, (,), {, }, [,], ^, ", ~, ?, :, /, \.

Procedure

1 With admin privileges, log in to NSX Manager.

2 On the homepage, enter a search pattern for an object or object type.

You can also select a recent query or a search query that you saved.

As you enter a search pattern, the search feature provides assistance by showing the applicable keywords.

Search	Search Query
Objects with Logical as the name or property	Logical
Exact logical switch name	display_name:LSP-301
Names with special characters such as, !	Logical\!

All the related search results are listed and grouped by resource type in different tabs.

You can click the tabs for specific search results for a resource type.

3 (Optional) In the search bar, click the save icon to save your refined search criteria.

4 In the search bar, click the  icon to open the advanced search column where you can refine your search.

5 Specify one or more criteria to refine your search.

- Name
- Resource Type
- Description
- ID
- Created by
- Modified by
- Tags
- Creation Date
- Modified Date

You can also view your recent search results and saved search criteria.

6 (Optional) Click **Clear All** to reset your advanced search criteria.

Filter by Object Attributes

When viewing objects in NSX Manager, you can filter the objects by one or more of their attributes. For example, when viewing details of Tier 0 gateways you can choose to filter by **Status** and view only those gateways that are **`Down`**.

The following types of filters are available:

- Predefined filters – A list of commonly used filters that you can apply to your objects.
- Text-based filter – A filter based on the attribute value that you enter. This filter is applicable only to the **Name**, **Tag**, **Path**, and **Description** attributes of the objects.
- Attribute-value pairs – An attribute drop-down menu that you can use to specify attribute-value pairs for filtering.

You can either use multiple attributes of an object or multiple values of a single attribute to filter objects. The AND operator is applied when you select multiple attributes whereas the OR operator is used when you specify multiple values of a single attribute.

Procedure

1. From your browser, log in with admin privileges to an NSX Manager at https://<nsx-manager-ip-address>.
2. Navigate to the tab that displays the objects you want to view.
3. Specify the attributes that you want to use to filter the objects.
   - Click  and select from a list of predefined filters.
   - Enter a value for the **Name**, **Tag**, **Path**, or **Description** attributes.
   - Select an attribute from the drop-down menu and specify its value. For example, **Status**: **`Down`**

   Objects satisfying your filter criteria are displayed.
4. (Optional) Click **Clear** to reset your filters.

Add a Compute Manager

A compute manager, for example, VMware vCenter, is an application that manages resources such as hosts and VMs.

NSX polls compute managers to collect cluster information from VMware vCenter.

For more information about VMware vCenter roles and privileges, see the *vSphere Security* document.

Prerequisites

- Verify that you use the supported vSphere version. See Supported vSphere version.

- IPv4 communication with VMware vCenter.
- Verify that you use the recommended number of compute managers. See https://configmax.vmware.com/home.
- Decide the hashing algorithm type you want to use for stamping NSX Manager thumbprint in compute manager extension. SHA1 and SHA256 algorithm types are supported. The default is SHA1. If you use SHA256 there might be communication issues between WCP component in VC and NSX Manager.
  - To set the hashing algorithm, run API `PUT https://<nsx-mgr>/api/v1/fabric/compute-managers/thumbprint-hashing-algorithm`

    ```
    {
        "hashing_algorithm_type": "SHA1"
    }
    ```

- Provide credentials of a VMware vCenter user. You can provide the credentials of VMware vCenter administrator, or create a role and a user specifically for NSX and provide this user's credentials. Add global permissions to the newly created user and role and select **Propogate to Children**.

  Create an admin role with the following VMware vCenter privileges:

Extension.Register extension
Extension.Unregister extension
Extension.Update extension
Sessions.Message
Sessions.Validate session
Sessions.View and stop sessions
Host.Configuration.Maintenance
Host.Configuration.NetworkConfiguration
Host.Local Operations.Create virtual machine
Host.Local Operations.Delete virtual machine
Host.Local Operations.Reconfigure virtual machine
Tasks
Scheduled task
Global.Cancel task
Permissions.Reassign role permissions
Resource.Assign vApp to resource pool
Resource.Assign virtual machine to resource pool
Virtual Machine.Configuration
Virtual Machine.Guest Operations

Virtual Machine.Provisioning
Virtual Machine.Inventory
Network.Assign network
vApp

To use the NSX license for the vSphere Distributed Switch 7.0 feature, the VMware vCenter user must either be an administrator, or the user must have *Global.Licenses* privileges and be a member of the *LicenseService.Administrators* group.

- Before you create a service account for the compute manager, add these additional VMware vCenter privileges to the admin user role:

Service Account Management.Administer
Permissions.Modify permission
Permissions.Modify role
VMware vSphere Lifecycle Manager.ESXi Health Perspectives.Read
VMware vSphere Lifecycle Manager.Lifecycle Manager: General Privileges.Read
VMware vSphere Lifecycle Manager.Lifecycle Manager: Image Privileges.Read
VMware vSphere Lifecycle Manager.Lifecycle Manager: Image Privileges.Write
VMware vSphere Lifecycle Manager.Lifecycle Manager: Image Remediation Privileges.Write
VMware vSphere Lifecycle Manager.Lifecycle Manager: Settings Privileges.Read
VMware vSphere Lifecycle Manager.Lifecycle Manager: Settings Privileges.Write
VMware vSphere Lifecycle Manager.Lifecycle Manager: General Privileges.Write

Procedure

1. From your browser, log in with admin privileges to an NSX Manager at https://<nsx-manager-ip-address>.
2. Select **System > Fabric > Compute Managers > Add Compute Manager**.
3. Complete the compute manager details.

Option	Description
Name and Description	Type the name to identify the VMware vCenter. You can optionally describe any special details such as, the number of clusters in the VMware vCenter.
Type	The default compute manager type is set to VMware vCenter.

Option	Description
Multi NSX	Starting with NSX 3.2.2, you can register the same vCenter Server with multiple NSX Managers. Enable this field if you want to allow multiple NSX instances to manage a single VMware vCenter. This functionality is supported from VMware vCenter 7.0 or later versions. **Note** Cannot be enabled on a Workload Control Plane (WCP) cluster or vSphere Lifecycle Manager (vLCM) cluster.
FQDN or IP Address	Type the FQDN or IP address of the VMware vCenter.
HTTPS Port of Reverse Proxy	The default port is 443. If you use another port, verify that the port is open on all the NSX Manager appliances. Set the reverse proxy port to register the compute manager in NSX.
Username and Password	Type the VMware vCenter login credentials.
SHA-256 Thumbprint	Type the VMware vCenter SHA-256 thumbprint algorithm value. If you configured the VMware vCenter WCP (Workload Control Plane) feature, using the SHA256 setting results in communication issues between the WCP component in VMware vCenter and NSX Manager. In such cases, use the SHA1 algorithm instead.
Create Service Account	Enable this field for features such as vSphere Lifecycle Manager that need to authenticate with NSX APIs. Log in with the administrator@vsphere.local credential to register a compute manager. After registration, the compute manager creates a service account. **Note** Service account creation is not supported on a global NSX Manager. If service account creation fails, the compute manager's registration status is set to `Registered with errors`. The compute manager is successfully registered. However, vSphere Lifecycle Manager cannot be enabled on NSX clusters. If a VMware vCenter admin deletes the service account after it was successfully created, vSphere Lifecycle Manager tries to authenticate the NSX APIs and the compute manager's registration status is set to `Registered with errors`.
Enable Trust	Enable this field to establish trust between NSX and compute manager, so that services running in vCenter Server can establish trusted communication with NSX. For example, for vSphere Lifecycle Manager to be enabled on NSX clusters, you must enable this field. Supported only on VMware vCenter 7.0 and later versions.
Access Level	Enable one of the options based on your requirement: ■ Full Access to NSX: Is selected by default. This access level gives the compute manager complete access to NSX. Full access ensures vSphere for Kubernetes and vSphere Lifecycle Manager can communicate with NSX. The VMware vCenter user's role must be set to an Enterprise Admin. ■ Limited Access to NSX: This access level ensures vSphere Lifecycle Manager can communicate with NSX. The VMware vCenter user's role must be set to Limited vSphere Admin.

If you left the thumbprint value blank, you are prompted to accept the server provided thumbprint.

After you accept the thumbprint, it takes a few seconds for NSX to discover and register the VMware vCenter resources.

Note If the FQDN, IP, or thumbprint of the compute manager changes after registration, edit the computer manager and enter the new values.

4 If the progress icon changes from **In progress** to **Not registered**, perform the following steps to resolve the error.

a Select the error message and click **Resolve**. One possible error message is the following:

```
Extension already registered at CM <vCenter Server name> with id <extension ID>
```

b Enter the VMware vCenter credentials and click **Resolve**.

If an existing registration exists, it will be replaced.

Results

It takes some time to register the compute manager with VMware vCenter and for the connection status to appear as `UP`.

You can click the compute manager's name to view the details, edit the compute manager, or to manage tags that apply to the compute manager.

After the VMware vCenter is successfully registered, do not power off and delete the NSX Manager VM without deleting the compute manager first. Otherwise, when you deploy a new NSX Manager, you will not be able to register the same VMware vCenter again. You will get the error that the VMware vCenter is already registered with another NSX Manager.

Note After a vCenter Server (VC) compute manager is successfully added, it cannot be removed if you successfully performed any of the following actions:

- Transport nodes are prepared using VDS that is dependent on the VC.
- Service VMs deployed on a host or a cluster in the VC using NSX service insertion.
- You use the NSX Manager UI to deploy Edge VMs or NSX Manager nodes on a host or a cluster in the VC.

If you try to perform any of these actions and you encounter an error (for example, installation failed), you can remove the VC if you have not successfully performed any of the actions listed above.

If you have successfully prepared any transport node using VDS that is dependent on the VC or deployed any VM, you can remove the VC after you have done the following:

- Unprepare all transport nodes. If uninstalling a transport node fails, you must force delete the transport node.
- Undeploy all service VMs, all NSX Edge VMs, and all NSX Manager nodes. The undeployment must be successful or in a failed state.
- If an NSX Manager cluster consists of nodes deployed from the VC (manual method) and nodes deployed from the NSX Manager UI, and you had to undeploy the manually deployed nodes, then you cannot remove the VC. To successfully remove the VC, ensure that you re-deploy an NSX Manager node from the VC.

This restriction applies to a fresh installation of NSX as well as an upgrade.

Replace Compute Manager

To replace a compute manager (VMware vCenter) with another compute manager (VMware vCenter), determine if the existing compute manager has NSX Manager cluster or NSX Edge nodes registered to it and third-party services running on the NSX cluster.

Before you can replace the VMware vCenter, you must delete any appliances or objects registered with the existing VMware vCenter and then deploy them on the new VMware vCenter.

Prerequisites

Procedure

1 Add a new compute manager. See Add a Compute Manager.

2 Replace NSX Manager cluster deployed from the **Add Appliance** wizard in the NSX Manager UI.

Note Do not follow this procedure to replace NSX Manager appliances deployed from VMware vCenter.

a Go to **System > Appliances > NSX Manager > Add NSX Appliance**.

b Verify whether the status of NSX Manager is stable. If not, ensure the status of NSX Manager is stable before proceeding with the next steps.

c Delete an existing NSX Manager on the old compute manager.

d Verify status of cluster by running the command, `get cluster status`.

e On the new compute manager, deploy a new NSX Manager.

f Once the new NSX Manager is fully operational and has joined the NSX Manager cluster, verify the cluster status.

g You can also view the status by running the command, `get cluster status`.

h Repeat these steps till all the existing NSX Manager are replaced by new ones and the cluster status is stable.

3 Add new NSX Edge nodes.

a To add Edge node, see Create an NSX Edge Transport Node.

b When creating a new NSX Edge node,

1 In the **Name and Descritpion** window, enter a new Host Name /FQDN.

2 In the **Configure Deployment** window, select a different compute manager.

c Verify NSX Edge node is deployed.

d Repeat the steps for all the NSX Edge nodes.

e All NSX Edge nodes are registered with the new compute manager.

f Go to **System > Fabric > Nodes > Edge Clusters**.

g Select the NSX Edge cluster and replace old NSX Edge nodes and add new NSX Edge nodes to the cluster. See Replacing an NSX Edge Transport Node in an NSX Edge Cluster.

h Click **Save** to save the NSX Edge cluster configuration.

4 Deploy third-party services.

Any third-party services registered with the existing VMware vCenter must be uninstalled from the old VMware vCenter and reinstalled on the new VMware vCenter.

a On the old compute manager, uninstall services introspecting east-west traffic. See Uninstall an East-West Traffic Introspection Service.

b On the old compute manager, uninstall services introspecting north-south traffic. See Uninstall a North-South Traffic Introspection Service.

c Install third-party services on the new compute manager.

See Deploy a Service for East-West Traffic Introspection.

See Deploy a Service for North-South Traffic Introspection.

Results

The old VMware vCenter is replaced by the new VMware vCenter and all appliances and objects are running on the new VMware vCenter.

Configuring Active Directory and Event Log Scraping

Active Directory is used in creating user-based Identity Firewall rules.

Windows 2008 is not supported as an Active Directory server or RDSH Server OS.

You can register one or more Windows domains with an NSX Manager. NSX Manager gets group and user information, and the relationship between them from each domain that it is registered. NSX Manager also retrieves Active Directory (AD) credentials.

Once the Active Directory is synced to the NSX Manager, you can create security groups based on user identity, and create identity-based firewall rules.

Scale limits for Active Directory, Event Log Scraping, and IDFW can be found on the VMware Configuration Maximums page.

Note For Identity Firewall rule enforcement, Windows Time service should be **on** for all VMs using Active Directory. This ensures that the date and time is synchronized between Active Directory and VMs. AD group membership changes, including enabling and deleting users, do not immediately take effect for logged in users. For changes to take effect, users must log out and then log back in. AD administrator's should force a logout when group membership is modified. This behavior is a limitation of Active Directory.

Prerequisites

If using event log scraping, make sure that NTP is configured correctly across all devices that will be using log scraping, for more information see Time Synchronization between NSX Manager, vIDM, and Related Components.

The domain account must have Active Directory read permission for all objects in the domain tree. The event log reader account must have read permissions for security event logs. See Enable Windows Security Log Access for the Event Log Reader.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Navigate to **System > Identity Firewall AD**.

3 Click **Add Active Directory**.

4 Enter the name of the active directory.

5 Enter the **NetBios Name** and **Base Distinguished Name**.

To retrieve the netBIOS name for your domain, enter `nbtstat -n` in a command window on a Windows Workstation that is part of a domain, or on a domain controller. In the NetBIOS Local Name Table, the entry with a <00> prefix and type Group is the NetBIOS name.

A base distinguished name (Base DN) is needed to add an Active Directory domain. A Base DN is the starting point that an LDAP server uses when searching for users authentication within an Active Directory domain. For example, if your domain name is corp.local the DN for the Base DN for Active Directory would be "DC=corp,DC=local".

6 Set the **Delta Synchronization Interval**, if necessary. A delta synchronization updates local AD objects that have changed since the last synchronization event.

Any changes made in Active Directory are NOT seen on NSX Manager until a delta or full synchronization has been performed.

7 Set the **LDAP Server**. See Add an LDAP Server for more information.

8 (Optional) Set the **Event Log Server**. Enter the Host IP or FQDN, user name and password, then click **Apply**.

9 Next to **Organization Units To Sync**, click **Sync all organization units and domains** or **Select organization units to sync**.

Groups that are moved out of the selected OrgUnits are not updated during a selective sync. Deleted groups are removed in a full sync, when all groups are updated.

Option	Description
Sync all organization units and domains	Full sync of all organization units is performed.
Select organization units to sync	Individually select organization units. If the parent is selected, the child units inside of the parent are automatically selected. You can also select all of the organization units by selecting the top **Organization Units** box, and then unselect the specific units you do not want to include in the sync. Only the selected organization units which are created and changed since the last delta sync will be updated during a selective sync.

10 Click **Save**.

11 The Active Directory screen appears, in a read only mode.

12 To edit an Active Directory:

a Click the three-dot menu () next to the Active Directory, and click **Edit**.

b You can now perform two actions: **Sync Delta**, or **Sync All**. For more information, see Synchronize Active Directory.

Enable Windows Security Log Access for the Event Log Reader

Read-only security access is used by event log scraper in IDFW.

After creating a new user account you must enable read-only security log access on a Windows 2008 and later server-based domain section, to grant the user read-only access

Members of the Event Log Readers group are granted permissions to read the event logs on the local computer. You must perform these steps on one Domain Controller of the domain, tree, or forest.

Prerequisites

The domain account must have Active Directory read permission for all objects in the domain tree. The event log reader account must have read permissions for security event logs.

Procedure

1 Navigate to **Start > Administrative Tools > Active Directory Users and Computers > .**

2 In the navigation tree, expand the node that corresponds to the domain for which you and to enable security log access.

3 Under the expanded node, select the **Builtin** node.

4 Double-click **Event Log Readers** in the list of groups.

5 Select the **Members** tab in the Event Log Readers Properties dialog box.

6 Click **Add**, and select the user or group you want to add to the Event Log Readers Group.

The **Select Users, Contacts, Computers, or Groups** dialog appears.

7 Click **OK** to close all open dialog boxes.

Add an LDAP Server

LDAP (Lightweight Directory Access Protocol) server configuration and functionality is only for use with Identity Firewall. LDAP provides a central place for authentication, meaning that when you configure a connection to your LDAP server, the user records are stored in your external LDAP server.

Prerequisites

The domain account must have AD read permission for all objects in the domain tree.

When there is a cluster of NSX Managers, all nodes need to be able to reach the LDAP server.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Navigate to **System > Identity Firewall AD**.

3 Click LDAP Server **Set**.

4 Click **Add LDAP Server**.

5 Enter the **Host** name of the LDAP server.

6 Select the active directory the LDAP server is connected to from the **Connected to (Directory)** drop-down menu.

7 (Optional) Select the **protocol**: LDAP (unsecured) or LDAPS (secured).

8 If LDAPS was selected, select the SHA-256 Thumbprint suggested by NSX Manager, or enter a SHA-256 Thumbprint.

9 Enter the **Port** number of the LDAP server.

   For local domain controllers, the default LDAP port 389 and LDAPS port 636 are used for the Active Directory sync, and should not be edited from the default values.

10 Enter the **username** and **password** of an Active Directory account with a minimum of read-only access to the Active Directory domain.

11 Click **ADD** and **APPLY**.

12 To verify that you can connect to the LDAP server, click **Test Connection**.

Synchronize Active Directory

Active Directory objects can be used to create security groups based on user identity, and identity-based firewall rules.

Note Do not enable Distributed Intrusion Detection Service (IDS) in an environment that is using Distributed Load Balancer. NSX does not support using IDS with a Distributed Load Balancer.

You can register an entire AD (Active Directory) domain to be used by IDFW (Identity Firewall), or you can synchronize a subset of a large domain. Once a domain is registered, NSX synchronizes all AD data required by IDFW. Selective sync is used for large active directory domains.

Selective sync allows you to selectively choose organizational units so that you do not have to sync the entire domain. Only the selected organization units which are created and changed since the last delta sync will be updated during a selective sync. Groups that are moved out of the selected organization units are not updated during a selective sync. Configuration maximums still apply selective sync. Deleted groups are removed in a full sync, when all groups are updated. To specify organization units for synchronization, see Configuring Active Directory and Event Log Scraping.

If you use the API to manually end a full sync after it is has begun, the sync stats will not be updated correctly.

Scale limits for Active Directory and IDFW can be found on the VMware Configuration Maximums page.

Note IDFW relies on the security and integrity of the guest operating system. There are multiple methods for a malicious local administrator to spoof their identity to bypass firewall rules. User identity information is provided by the Guest Introspection Agent inside guest VMs. Security administrators must ensure that NSX Guest Introspection Agent is installed and running in each guest VM. Logged-in users should not have the privilege to remove or stop the agent.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Navigate to **System > Identity Firewall AD**.

3 Click the three button menu ( ) next to the Active Directory that you want to synchronize, and select one of the following:

Option	Description
Sync All	Full sync of all data is performed from the Active Directory, regardless of the state of sync on NSX.
Sync Delta	Perform a delta synchronization, where local AD objects that have changed since the last synchronization are updated. A full sync of all data is not performed. Deleted groups are removed during Sync All, when all groups are updated.

4 Click **Save**.

5 Click **View Sync Status** to see the current state of the Active Directory, the previous synchronization state, the synchronization status, and the last synchronization time.

Remove NSX Extension from VMware vCenter

When you add a compute manager, NSX Manager adds its identity as an extension in VMware vCenter. If you remove the compute manager, the extension in VMware vCenter will be removed automatically. If the extension is not removed for some reason, you can maually remove the extension with the following procedure.

Prerequisites

Enable access to the VMware vCenter Managed Object Browser (MOB) by following the procedure in https://kb.vmware.com/s/article/2042554.

Procedure

1. Login to the MOB at `https://<vCenter Server hostname or IP address>/mob`.
2. Click the **content** link, which is the value for the **content** property in the Properties table.
3. Click the **ExtensionManager** link, which is the value for **extensionManager** property in the Properties table.
4. Click the **UnregisterExtension** link in the Methods table.
5. Enter **`com.vmware.nsx.management.nsxt`** in the **value** text field.
6. Click the **Invoke Method** link on the right hand side of the page below the Parameters table.

   The method result says `void` but the extension will be removed.
7. To make sure the extension is removed, click the **FindExtension** method on the previous page and invoke it by entering the same value for the extension.

   The result should be `void`.

Managing the NSX Manager Cluster

You can reboot an NSX Manager if it becomes inoperable. You can also change the IP address of an NSX Manager.

In a production environment, it is highly recommended that the NSX Manager cluster has three members to provide high availability. If you delete an NSX Manager and deploy a new one, the new NSX Manager can have the same or a different IP address.

Note The primary NSX Manager node is the node that you create first, before you create a manager cluster. This node cannot be deleted. After you deploy two more manager nodes from the primary manager node's UI to form a cluster, only the second and the third manager nodes have the option (from the gear icon) to be deleted. For information about removing and adding a manager node, see Change the IP Address of an NSX Manager.

View the Configuration and Status of the NSX Manager Cluster

You can view the configuration and status of the NSX Manager cluster from the NSX Manager UI. You can get additional information using the CLI.

Procedure

1. From your browser, log in with admin privileges to an NSX Manager at https://*nsx-manager-ip-address*.

2 Select **System > Overview**

The status of the NSX Manager cluster is displayed.

3 To see additional information about the configuration, run the following CLI command:

```
manager1> get cluster config
Cluster Id: 18807edd-56d1-4107-b7b7-508d766a08e3
Cluster Configuration Version: 3
Number of nodes in the cluster: 3

Node UUID: 43cd0642-275c-af1d-fe46-1f5200f9e5f9
Node Status: JOINED
    ENTITY                              UUID                                    IP
ADDRESS      PORT     FQDN
    HTTPS                               5c8d01f1-f3ee-4f94-b517-a093d8fbfad3
10.160.71.225   443      ychin-nsxmanager-ob-12065118-1-F5
    CONTROLLER                          06fd0574-69c0-432e-a8af-53d140dbef8f
10.160.71.225   -        ychin-nsxmanager-ob-12065118-1-F5
    CLUSTER_BOOT_MANAGER                da8d535e-7a0c-4dd8-8919-d88bdde006b8
10.160.71.225   -        ychin-nsxmanager-ob-12065118-1-F5
    DATASTORE                           3c9c4ec1-afef-47bd-aadb-1ed6a5536bc4
10.160.71.225   9000     ychin-nsxmanager-ob-12065118-1-F5
    MANAGER                             eb5e8922-23bd-4c3a-ae22-d13d9195a6bc
10.160.71.225   -        ychin-nsxmanager-ob-12065118-1-F5
    POLICY                              f9da1039-08ad-4a20-bacc-5b91c5d67730
10.160.71.225   -        ychin-nsxmanager-ob-12065118-1-F5

Node UUID: 8ebb0642-201e-6a5f-dd47-a1e38542e672
Node Status: JOINED
    ENTITY                              UUID                                    IP
ADDRESS      PORT     FQDN
    HTTPS                               3757f155-8a5d-4b53-828f-d67041d5a210
10.160.93.240   443      ychin-nsxmanager-ob-12065118-2-F5
    CONTROLLER                          7b1c9952-8738-4900-b68b-ca862aa4f6a9
10.160.93.240   -        ychin-nsxmanager-ob-12065118-2-F5
    CLUSTER_BOOT_MANAGER                b5e12db1-5e0d-4e33-a571-6ba258dceb2e
10.160.93.240   -        ychin-nsxmanager-ob-12065118-2-F5
    DATASTORE                           bee1f629-4e23-4ab8-8083-9e0f0bb83178
10.160.93.240   9000     ychin-nsxmanager-ob-12065118-2-F5
    MANAGER                             45ccd6e3-1497-4334-944c-e6bbcd5c723e
10.160.93.240   -        ychin-nsxmanager-ob-12065118-2-F5
    POLICY                              d5ba5803-b059-4fbc-897c-3aace8cf1219
10.160.93.240   -        ychin-nsxmanager-ob-12065118-2-F5

Node UUID: 2e7e0642-df4a-b2ec-b9e8-633d1469f1ea
Node Status: JOINED
    ENTITY                              UUID                                    IP
ADDRESS      PORT     FQDN
    HTTPS                               bce3cc4c-7d60-45e2-aa7b-cdc75e445a14
10.160.76.33    443      ychin-nsxmanager-ob-12065118-3-F5
    CONTROLLER                          ced46f5c-9e52-4b31-a1cb-b3dead991c71
10.160.76.33    -        ychin-nsxmanager-ob-12065118-3-F5
    CLUSTER_BOOT_MANAGER                88b70d31-3428-4ccc-ab57-55859f45030c
10.160.76.33    -        ychin-nsxmanager-ob-12065118-3-F5
```

```
    DATASTORE                           fb4aec3c-cae3-4386-b5b9-c0b99b7d9048
10.160.76.33    9000     ychin-nsxmanager-ob-12065118-3-F5
    MANAGER                             82b07440-3ff6-4f67-a1c9-e9327d1686ad
10.160.76.33    -        ychin-nsxmanager-ob-12065118-3-F5
    POLICY                              61f21a78-a56c-4af1-867b-3f24132d53c7
10.160.76.33    -        ychin-nsxmanager-ob-12065118-3-F5
```

4 To see additional information about the status, run the following CLI command:

```
manager1> get cluster status
Cluster Id: 18807edd-56d1-4107-b7b7-508d766a08e3
Group Type: DATASTORE
Group Status: STABLE

Members:
    UUID                                        FQDN
IP              STATUS
    43cd0642-275c-af1d-fe46-1f5200f9e5f9        ychin-nsxmanager-ob-12065118-1-F5
10.160.71.225    UP
    8ebb0642-201e-6a5f-dd47-a1e38542e672        ychin-nsxmanager-ob-12065118-2-F5
10.160.93.240    UP
    2e7e0642-df4a-b2ec-b9e8-633d1469f1ea        ychin-nsxmanager-ob-12065118-3-F5
10.160.76.33     UP

Group Type: CLUSTER_BOOT_MANAGER
Group Status: STABLE

Members:
    UUID                                        FQDN
IP              STATUS
    43cd0642-275c-af1d-fe46-1f5200f9e5f9        ychin-nsxmanager-ob-12065118-1-F5
10.160.71.225    UP
    8ebb0642-201e-6a5f-dd47-a1e38542e672        ychin-nsxmanager-ob-12065118-2-F5
10.160.93.240    UP
    2e7e0642-df4a-b2ec-b9e8-633d1469f1ea        ychin-nsxmanager-ob-12065118-3-F5
10.160.76.33     UP

Group Type: CONTROLLER
Group Status: STABLE

Members:
    UUID                                        FQDN
IP              STATUS
    7b1c9952-8738-4900-b68b-ca862aa4f6a9        ychin-nsxmanager-ob-12065118-2-F5
10.160.93.240    UP
    ced46f5c-9e52-4b31-a1cb-b3dead991c71        ychin-nsxmanager-ob-12065118-3-F5
10.160.76.33     UP
    06fd0574-69c0-432e-a8af-53d140dbef8f        ychin-nsxmanager-ob-12065118-1-F5
10.160.71.225    UP

Group Type: MANAGER
Group Status: STABLE

Members:
    UUID                                        FQDN
```

```
IP               STATUS
    43cd0642-275c-af1d-fe46-1f5200f9e5f9       ychin-nsxmanager-ob-12065118-1-F5
10.160.71.225    UP
    8ebb0642-201e-6a5f-dd47-a1e38542e672       ychin-nsxmanager-ob-12065118-2-F5
10.160.93.240    UP
    2e7e0642-df4a-b2ec-b9e8-633d1469f1ea       ychin-nsxmanager-ob-12065118-3-F5
10.160.76.33     UP

Group Type: POLICY
Group Status: STABLE

Members:
    UUID                                       FQDN
IP               STATUS
    43cd0642-275c-af1d-fe46-1f5200f9e5f9       ychin-nsxmanager-ob-12065118-1-F5
10.160.71.225    UP
    8ebb0642-201e-6a5f-dd47-a1e38542e672       ychin-nsxmanager-ob-12065118-2-F5
10.160.93.240    UP
    2e7e0642-df4a-b2ec-b9e8-633d1469f1ea       ychin-nsxmanager-ob-12065118-3-F5
10.160.76.33     UP

Group Type: HTTPS
Group Status: STABLE

Members:
    UUID                                       FQDN
IP               STATUS
    43cd0642-275c-af1d-fe46-1f5200f9e5f9       ychin-nsxmanager-ob-12065118-1-F5
10.160.71.225    UP
    8ebb0642-201e-6a5f-dd47-a1e38542e672       ychin-nsxmanager-ob-12065118-2-F5
10.160.93.240    UP
    2e7e0642-df4a-b2ec-b9e8-633d1469f1ea       ychin-nsxmanager-ob-12065118-3-F5
10.160.76.33     UP
```

Update API Service Configuration of the NSX Manager Cluster

You can modify the API service properties of the NSX Manager cluster, such as TLS protocol version, cipher suites, and so on.

The following procedure explains the workflow of running the NSX API service calls to disable the TLS 1.1 protocol, and to enable or disable the cipher suites in the API service configuration.

For a detailed information about the API schema, example request, example response, and error messages of the NSX API service, you must read the *NSX API Guide*.

Procedure

1 Run the following GET API to read the configuration of the NSX API service:

```
GET https://<NSX-Manager-IP>/api/v1/cluster/api-service
```

The API response contains the list of cipher suites and TLS protocols.

2 Disable the TLS 1.1 protocol.

   a Set `TLSv1.1` to `enabled = false`.

   b Run the following PUT API to send the changes to the NSX API server:

   ```
   PUT https://<NSX-Manager-IP>/api/v1/cluster/api-service
   ```

3 Enable or disable the cipher suites.

   a Set one or more cipher names to `enabled = false` or `enabled = true` depending on your requirement.

   b Run the following PUT API to send the changes to the NSX API server:

   ```
   PUT https://<NSX-Manager-IP>/api/v1/cluster/api-service
   ```

Results

The API service on each NSX Manager node restarts after it is updated using the API. There might be a delay of up to a minute between the time the API call completes and when the new configuration comes into effect. The changes in the API service configuration are applied to all the nodes in the NSX Manager cluster.

Shut Down and Power On the NSX Manager Cluster

If you need to shut down the NSX Manager cluster, use the following procedure.

Procedure

1 To shut down an NSX Manager cluster, shut down one manager node at a time. You can log in to the command-line interface (CLI) of a manager node as `admin` and run the command `shutdown`, or shut down the manager node VM from VMware vCenter.

   Make sure that the VM is powered off in VMware vCenter before proceeding to the next one.

2 To power on an NSX Manager cluster, power on one manager node VM at a time in VMware vCenter.

   Make sure the node health API returns `healthy: true` before you proceed to the next node.

```
curl -u admin:${PASSWORD} -i -k https://$IP/api/v1/reverse-proxy/node/health
```

```
{
"healthy" : true,
"components_health" : "SEARCH:UP, MANAGER:UP, UI:UP, NODE_MGMT:UP"
}
```

Reboot an NSX Manager

You can reboot an NSX Manager with a CLI command to recover from critical errors.

If you need to reboot multiple NSX Managers, you must reboot them one at a time. Wait for the rebooted NSX Manager to be online before rebooting another.

Procedure

1 Log in to the CLI of the NSX Manager.

2 Run the following command.

```
nsx-manager> reboot
Are you sure you want to reboot (yes/no): y
```

Change the IP Address of an NSX Manager

You can change the IP address of an NSX Manager in an NSX Manager cluster. This section describes several approaches.

For example, if you have a cluster consisting of Manager A, Manager B, and Manager C, you can change the IP address of one or more of the managers in the following ways:

- Scenario A:
  - Manager A has IP address 172.16.1.11.
  - Manager B has IP address 172.16.1.12.
  - Manager C has IP address 172.16.1.13.
  - Add Manager D with a new IP address, for example, 192.168.55.11.
  - Remove Manager A.
  - Add Manager E with a new IP address, for example, 192.168.55.12.
  - Remove Manager B.
  - Add Manager F with a new IP address, for example, 192.168.55.13.
  - Remove Manager C.
- Scenario B:
  - Manager A has IP address 172.16.1.11.
  - Manager B has IP address 172.16.1.12.
  - Manager C has IP address 172.16.1.13.
  - Add Manager D with a new IP address, for example, 192.168.55.11.
  - Add Manager E with a new IP address, for example, 192.168.55.12.
  - Add Manager F with a new IP address, for example, 192.168.55.13.
  - Remove Manager A, Manager B, and Manager C.
- Scenario C:
  - Manager A has IP address 172.16.1.11.

- Manager B has IP address 172.16.1.12.
- Manager C has IP address 172.16.1.13.
- Remove Manager A.
- Add Manager D with a new IP address, for example, 192.168.55.11.
- Remove Manager B.
- Add Manager E with a new IP address, for example, 192.168.55.12.
- Remove Manager C.
- Add Manager F with a new IP address, for example, 192.168.55.13.

The first two scenarios require additional virtual RAM, CPU and disk for the additional NSX Managers during this IP address change.

Scenario C is not recommended because it temporarily reduces the number of NSX Managers and a loss of one of the two active managers during the IP address change will have an impact on the operations of NSX. This scenario is for a situation where additional virtual RAM, CPU and disk are not available and an IP address change is required.

Note If you are using the cluster VIP feature, you must either use the same subnet for the new IP addresses or disable the cluster VIP during the IP address changes because the cluster VIP requires all NSX Managers to be in the same subnet.

Prerequisites

Familiarize yourself with how to deploy an NSX Manager into a cluster. For more information, see the *NSX Installation Guide*.

Procedure

1 If the NSX Manager you want to remove was deployed manually, perform the following steps.

   a Run the following CLI command to detach the NSX Manager from the cluster.

   ```
   detach node <node-id>
   ```

   b Delete the NSX Manager VM.

2 If the NSX Manager you want to delete was deployed automatically through the NSX Manager UI, perform the following steps.

   a From your browser, log in with administrator privileges to an NSX Manager at https://*nsx-manager-ip-address*.

   This NSX Manager must not be the one that you want to delete.

   b From the **Systems** tab, click **NSX Management Nodes**.

   The status of the NSX Manager cluster is displayed.

   c For the NSX Manager that you want to delete, click the gear icon and select **Delete**.

3 Deploy a new NSX Manager.

Resize an NSX Manager Node

There are two ways to change the memory and CPU resources of an NSX Manager node in a cluster.

Note that in normal operating conditions all three manager nodes must have the same CPU and memory resources. A mismatch of CPU or memory between NSX Managers in an NSX Manager cluster should only be done when transitioning from one size of NSX Manager to another size.

If you have configured resource allocation reservation for the NSX Manager VMs in vCenter Server, you might need to adjust the reservation. For more information, see the vSphere documentation.

Option 1 (resize a manager node with the same IP address) requires less effort. NSX requires that two managers are available at all times. If you have cluster VIP (virtual IP) configured, there will be a brief outage when the VIP switches to another node in the cluster. You can access the other two nodes directly during the outage if the VIP-assigned node is shut down for resizing. If you have deployed a load balancer for the manager nodes, health checks will be triggered when a manager goes offline. The load balancer should direct traffic to another node. Choose this option if you do not want to change the IP address of the manager nodes.

For option 2 (resize a manager node with a different IP address), you will need IP addresses for the three new managers. If you have cluster VIP configured, there will be a brief outage when the VIP switches to another node in the cluster. You can access the other two nodes directly during the outage in case the VIP-assigned node is deleted. If you have deployed a load balancer for the manager nodes, health checks will be triggered when a manager goes offline. The load balancer should direct traffic to another node. After all the steps are completed, you will need to reconfigure the load balancer (add the new managers and remove the old managers).

Prerequisites

- Verify that the new size satisfies the system requirements for a manager node. For more information, see "NSX Manager VM and Host Transport Node System Requirements" in the *NSX Installation Guide*.
- Familiarize yourself with how to run CLI commands. For more information, see the *NSX Command-Line Interface Reference*. Also familiarize yourself with how to change the memory and CPU resources of a VM. For more information, see the vSphere documentation.
- Familiarize yourself with the requirements of an NSX Manager cluster. For more information, see "NSX Manager Cluster Requirements" in the *NSX Installation Guide*.
- Familiarize yourself with how to deploy an NSX Manager into a cluster. For more information, see "Deploy NSX Manager Nodes to Form a Cluster from the UI" in the *NSX Installation Guide*.

Procedure

- Option 1: Resize a manage node with the same IP address

  Option 1a: Change the CPU and/or memory of the existing manager nodes. You must make the change to one manager at a time so that two managers are available at all times.

  a Log in to a manager's CLI as **admin** and run the `shutdown` command.

  b From NSX Manager UI, verify that the state of the manager cluster is DEGRADED.

  c From vSphere, change the memory and/or CPU resources of the manager VM that was shut down.

  d From vSphere, power on the VM. From NSX Manager UI, wait for the state of the manager cluster to be STABLE.

  e Repeat steps 1 to 4 for the other two manager VMs.

  Option 1b: Deploy new manager nodes.

  a From NSX Manager UI, delete a manager node that was deployed from NSX Manager UI.

  b From NSX Manager UI, deploy a new manager node with the new size into the cluster with an IP address that is the same as the one used by the manager node that was deleted in step 1.

  c From NSX Manager UI, wait for the state of the manager cluster to be STABLE.

  d Repeat steps 1 to 3 for the other manager node that was deployed from NSX Manager UI.

  e For the manually-deployed manager node, log in to its CLI as **admin** and run the `shutdown` command.

  f From another manager node, log in to its CLI as **admin** and run the `get cluster config` command to get the node ID of the manually-deployed manager node. Then run the `detach node <node-id>` command to detach the manually-deployed manager node from the cluster.

  g From vSphere, delete the manually-deployed manager node VM.

  h From NSX Manager UI, deploy a new manager node with the new size into the cluster with an IP address that is the same as the one used by the manually-deployed manager node.

  i From NSX Manager UI, wait for the state of the manager cluster to be STABLE.

- Option 2: Resize a manage node with a different IP address

  a If you have VIP configured and if the new addresses and old addresses are in different subnets, from NSX Manager UI, remove the VIP.

  You must access NSX Manager using the manager's IP address and not the VIP address.

  b From NSX Manager UI, deploy a new manager node with the new size into the cluster with an IP address that is different from the ones used by the current manager nodes.

c From NSX Manager UI, verify that the state of the manager cluster is STABLE.

d From NSX Manager UI, delete an old manager node that was deployed from NSX Manager UI.

e Repeat steps 1 to 3 for the other manager node that was deployed from NSX Manager UI.

f For the manually-deployed manager node, log in to its CLI as **admin** and run the `shutdown` command.

g From another manager node, log in to its CLI as **admin** and run the `get cluster config` command to get the node ID of the manually-deployed manager node. Then run the `detach node <node-id>` command to detach the manually-deployed manager node from the cluster.

h From vSphere, delete the manually-deployed manager node VM.

i From NSX Manager UI, deploy a new manager node with the new size into the cluster with an IP address that is different from the one used by the manually-deployed manager node.

j From NSX Manager UI, wait for the state of the manager cluster to be STABLE.

k If you removed the old VIP in step 1, from NSX Manager UI, configure a new VIP. It must be in the same subnet that the new IP addresses are in.

Replacing an NSX Edge Transport Node in an NSX Edge Cluster

You can replace an NSX Edge transport node in an NSX Edge cluster using the NSX Manager UI or the API.

Important NSX Edge cluster does not support heterogenous NSX Edge nodes, such as a cluster formed by NSX Edge VMs and Bare Metal NSX Edge nodes. However, you can use heterogenous NSX Edge nodes only when you want to migrate these nodes.

Replace an NSX Edge Transport Node Using the NSX Manager UI

The following procedure describes replacing an NSX Edge transport node in an NSX Edge cluster that has both VM NSX Edge and Bare Metal NSX Edge transport nodes using the NSX Manager UI. You can replace an VM NSX Edge with Bare Metal NSX Edge or the other way around. You can replace the Edge transport node regardless of whether it is running or not.

If the NSX Edge node to be replaced is not running, the new NSX Edge node can have the same management IP address and TEP IP address. If the NSX Edge node to be replaced is running, the new NSX Edge node must have a different management IP address and TEP IP address.

Prerequisites

- Familiarize yourself with the procedure to install an NSX Edge node, join the Edge node with the management plane, and create an NSX Edge transport node. For more information, see the *NSX Installation Guide*.
- Both VM NSX Edge and Bare Metal NSX Edge transport nodes must have the same VLAN connectivity to the physical Top of Rack (TOR) switches.

Procedure

1. If you want the new NSX Edge transport node to have the same configurations as the NSX Edge transport node to be replaced, make the following API call to find the configurations:

   `GET https://<nsx-manager-IP>/api/v1/transport-nodes/<tn-id>`

2. Follow the procedures in the *NSX Installation Guide* to install and configure an Edge transport node.

   If you want this NSX Edge transport node to have the same configurations as the NSX Edge transport node to be replaced, use the configurations obtained in step 1.

3. In NSX Manager, select **System > Fabric > Nodes > Edge Clusters**.
4. Select an NSX Edge cluster by clicking the checkbox in the first column.
5. SSH into the NSX Edge nodes where Tier-0 is hosted.
6. Run `get logical router`. Check the VRF ID for Tier-0 Service Router (SR) on all NSX Edge nodes in the NSX Edge Cluster.
7. If the VRF id of the Tier-0 SR is 1, run `vrf 1`.
8. To check the output of the service router, run `get high-availability status`.
9. Enable maintenance mode on one of NSX Edge nodes that has Tier-0 SR in Standby. In the Edge CLI console, run `set maintenance-mode enabled`.

   This NSX Edge node could have Tier-1 SRs in Active state. Putting NSX Edge node into maintenance mode triggers a HA failover and all Tier-1 or Tier-0 SRs on this NSX Edge Node go into Standby state on this NSX Edge node. This might cause traffic disruption for the active SRs on this NSX Edge node because of failover of Tier-1 or Tier-0 SRs.

10. Ensure that Bare Metal NSX Edge transport node is not a part of any other cluster.
11. Click **Actions > Replace Edge Cluster Member**.

    It is recommended that you place the transport node being replaced in maintenance mode. If the transport node is not running, you can safely ignore this recommendation.

12. Select the VM NSX Edge transport node to be replaced from the dropdown list.
13. Select the Bare Metal NSX Edge transport node replacement node from the dropdown list.
14. Click **Save**.

15 Verify that the Bare Metal NSX Edge transport node has moved into existing Edge VM Cluster.

16 To verify that Tier-0 and Tier-1 gateways have moved from NSX Edge VM (in maintenance mode) to Bare Metal NSX Edge transport node, run `get logical router`.

17 Repeat the previous steps to move another VM NSX Edge with Bare Metal NSX Edge transport node.

18 Verify E-W and N-S connectivity from the workloads connected to Tier-1 or Tier-0 LRs.

Results

If you are running an NSX version earlier than 3.1.3, after replacing the NSX Edge transport node, you might see the alarm `All BGP/BFD sessions are down`. To resolve the issue, follow the workaround instructions in the KB article https://kb.vmware.com/s/article/83983.

What to do next

Replacing a VM NSX Edge VM with Bare Metal NSX Edge node does not automatically rebalance the Tier-1 gateways across NSX Edge nodes. You need to manually reconfigure each Tier-1 gateway.

Replace an NSX Edge Transport Node Using the API

The following procedure describes replacing an NSX Edge transport node in an NSX Edge cluster using the NSX API. You can replace the Edge transport node regardless of whether it is running or not.

If the Edge node to be replaced is not running, the new Edge node can have the same management IP address and TEP IP address. If the Edge node to be replaced is running, the new Edge node must have a different management IP address and TEP IP address.

Prerequisites

- Familiarize yourself with the procedure to install an NSX Edge node, join the Edge node with the management plane, and create an NSX Edge transport node. For more information, see the *NSX Installation Guide*.
- Ensure the new NSX Edge node that will use to replace the old NSX Edge node is ready. For more information, see the *NSX Installation Guide*.

Procedure

1 If you want the new NSX Edge transport node to have the same configurations as the old NSX Edge transport node to be replaced, make the following API call to find the configurations:

`GET https://<nsx-manager-IP>/policy/api/v1/infra/sites/<site-id>/enforcement-points/<enforcementpoint-id>/host-transport-nodes`


```
{
    "node_id": "250175b8-223b-11ed-826e-b07b25e93f64",
    "host_switch_spec": {
        "host_switches": [
            {
                "host_switch_name": "nsxHostSwitch",
                "host_switch_id": "809299a2-c090-4543-8747-d200e12cd2ea",
                "host_switch_type": "NVDS",
                "host_switch_mode": "STANDARD",
                "host_switch_profile_ids": [
                    {
                        "key": "UplinkHostSwitchProfile",
                        "value": "57da58fa-bce6-448b-8db3-874ceff59656"
                    },
                    {
                        "key": "LldpHostSwitchProfile",
                        "value": "9e0b4d2d-d155-4b4b-8947-fbfe5b79f7cb"
                    }
                ],
                "pnics": [
                    {
                        "device_name": "fp-eth0",
                        "uplink_name": "Uplink1"
                    },
                    {
                        "device_name": "fp-eth2",
                        "uplink_name": "Uplink2"
                    },
                    {
                        "device_name": "fp-eth4",
                        "uplink_name": "Uplink3"
                    },
                    {
                        "device_name": "fp-eth6",
                        "uplink_name": "Uplink4"
                    }
                ],
                "is_migrate_pnics": false,
                "ip_assignment_spec": {
                    "ip_pool_id": "82f8ae96-992b-45c6-8376-777b82bfeb1d",
                    "resource_type": "StaticIpPoolSpec"
                },
                "cpu_config": [],
                "transport_zone_endpoints": [
                    {
```

```
                        "transport_zone_id": "15897bda-802f-4481-b9fd-4e5cc1ef084b",
                        "transport_zone_profile_ids": [
                            {
                                "resource_type": "BfdHealthMonitoringProfile",
                                "profile_id": "52035bb3-ab02-4a08-9884-18631312e50a"
                            }
                        ]
                    },
                    {
                        "transport_zone_id": "4a237a28-050e-4499-a241-0eb0c9dad97f",
                        "transport_zone_profile_ids": [
                            {
                                "resource_type": "BfdHealthMonitoringProfile",
                                "profile_id": "52035bb3-ab02-4a08-9884-18631312e50a"
                            }
                        ]
                    }
                ],
                "pnics_uninstall_migration": [],
                "vmk_uninstall_migration": [],
                "not_ready": false
            }
        ],
        "resource_type": "StandardHostSwitchSpec"
    },
    "maintenance_mode": "DISABLED",
    "node_deployment_info": {
        "deployment_type": "PHYSICAL_MACHINE",
        "node_settings": {
            "hostname": "w1-hs2-m2716.eng.vmware.com",
            "enable_ssh": true,
            "allow_ssh_root_login": false
        },
        "resource_type": "EdgeNode",
        "external_id": "250175b8-223b-11ed-826e-b07b25e93f64",
        "ip_addresses": [
            "10.196.145.177"
        ],
        "id": "250175b8-223b-11ed-826e-b07b25e93f64",
        "display_name": "w1-hs2-m2716.eng.vmware.com",
        "tags": [],
        "_revision": 3
    },
    "is_overridden": false,
    "failure_domain_id": "4fc1e3b0-1cd4-4339-86c8-f76baddbaafb",
    "resource_type": "TransportNode",
    "id": "250175b8-223b-11ed-826e-b07b25e93f64",
    "display_name": "w1-hs2-m2716.eng.vmware.com",
    "tags": [],
    "_create_time": 1661187299037,
    "_create_user": "admin",
    "_last_modified_time": 1661255498968,
    "_last_modified_user": "admin",
```

```
    "_system_owned": false,
    "_protection": "NOT_PROTECTED",
    "_revision": 3
}
```

2 Note the transport node ID of the node to be replaced "55120a1a-51c6-4c20-b4a3-6f59662c9f6a".

3 Prepare a new NSX Edge Transport Node that will replace the old NSX Edge node. See the *Create an NSX Edge Transport Node* in the *NSX Installation Guide*.

Note the following configuration points while preparing the new NSX Edge node:

- Do not use the same IP address of the old NSX Edge node if it is running, `"ip_addresses": ["10.161.68.92"]`.
- Do not use the same TEP IP address if the old NSX Edge node is running

4 Make an API call to retrieve the member index of the transport node that has to be replaced.

`GET https://<nsx-manager-IP>/api/v1/edge-clusters`

```
....
    {
      "resource_type": "EdgeCluster",
      "description": "",
      "id": "9a302df7-0833-4237-af1f-4d826c25ad78",
      "display_name": "Edge-Cluster-1",
...
      "members": [
        {
          "member_index": 0,
          "transport_node_id": "55120a1a-51c6-4c20-b4a3-6f59662c9f6a"
        },
        {
          "member_index": 1,
          "transport_node_id": "890f0e3c-aa81-46aa-843b-8ac25fe30bd3"
        }
      ],
```

5 Make an API call to replace a transport node in an NSX Edge cluster. The `member_index` must match the index of the transport node to be replaced.

For example, the transport node `TN-edgenode-01a` (73cb00c9-70d0-4808-abfe-a12a43251133) has failed and is replaced by transport node `TN-edgenode-03a` (890f0e3c-aa81-46aa-843b-8ac25fe30bd3) in NSX Edge cluster `Edge-Cluster-1` (9a302df7-0833-4237-af1f-4d826c25ad78).

```
POST http://<nsx-manager-IP>/api/v1/edge-clusters/9a302df7-0833-4237-af1f-4d826c25ad78?
action=replace_transport_node
{
    "member_index": 0,
    "transport_node_id" : "890f0e3c-aa81-46aa-843b-8ac25fe30bd3"
}
```

Results

If you are running an NSX version earlier than 3.1.3, after replacing the NSX Edge transport node, you might see the alarm "All BGP/BFD sessions are down." To resolve the issue, follow the workaround instructions in the KB article https://kb.vmware.com/s/article/83983.

Managing Resource Reservations for an Edge VM Appliance

NSX uses vSphere resource allocation to reserve resources for NSX Edge appliances. You can tune the CPU and memory resources reserved for NSX Edge to ensure optimal use of resources on an NSX Edge.

For maximum performance NSX Edge VM appliance must be assigned 100% of the available resources. If you customize resources allocated to the NSX Edge VM, turn back the allocation later to 100% to get maximum performance.

For auto-deployed NSX Edge appliances, you can change the resource allocation from NSX Manager. However, if an NSX Edge appliance is deployed from vSphere, you can only manage resource reservations for that NSX Edge VM from vSphere.

As per the resource requirements of the Edge VM deployed in your environment, there are two ways to manage reservations:

- Default values assigned to give 100% resource reservations.
- Custom values assigned to give 0–100% resource reservations.

Default Reservations

Assumes the NSX Edge set to the **High** priority. The level of priority importance defines the number of vCPU shares and memory assigned to the NSX Edge. To assign custom values, you can change the relative priority assigned to the NSX Edge.

Resource constraints for different form factors set with Normal priority:

Form Factor	Number of vCPUs	vCPU Shares	RAM (GB)
Small	2	2000	4
Medium	4	4000	8
Large	8	8000	32
XLarge	16	16000	64

You can tune reservations of an NSX Edge appliance by considering two parameters:

- Relative priority assigned to a VM
- Pre-assigned resource constraints for a VM form factor

Custom Reservations

Assign relative priority for an NSX Edge appliance. You can change the relative importance of an NSX Edge appliance to assign the following resource requirements:

Relative Importance	CPU Shares (shares per vCPU)	Memory (shares per MB configured virtual machine memory)
Extra High	4000	40
High	2000	20
Normal	1000	10
Low	500	5

For example, a High relative importance to an NSX Edge appliance deployed in a medium form factor assigns the following vCPU and memory shares:

- 4 (vCPUs) X 8000 (vCPU share value) = 32000 shares of vCPU
- 20 (GB RAM) X 1000 = 20000 shares of memory

Note Before assigning a CPU value in MHz to guarantee the allocated CPU cycles for an NSX EdgeVM, ensure that the relative importance is set to Low. If the relative importance is set to Normal or High with a custom CPU value in MHz, the VM deployment might face issues due to resource constraints.

Tune Resource Reservations for an NSX Edge Appliance

You can tune resource reservations on an NSX Edge VM appliance. By default, 100% resources are allocated to an NSX Edge VM. Flexibility to change resource reservations avoids the need to add additional capacity to the vCenter Server and the need to reduce current reservations on other non-Edge VMs.

Prerequisites

- Verify that the cluster has sufficient capacity to avoid failures.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Select **System > Fabric > Nodes > Edge Transport Nodes**.
3. Select the NSX Edge transport node.
4. Click **Actions > Change Edge VM Resource Reservations**.

5 In the **Change Edge VM Resource Reservations** window, you can customize the existing resource allocation applied to the Edge transport node.

Action	Description
CPU Reservation Priority	Low - 2000 shares Normal - 4000 shares High - 8000 shares Extra High - 10000 shares
Memory Reservation (%)	Reservation percentage is relative to the pre-defined value in the form factor. 100 indicates 100% of memory is reserved for the NSX Edge VM. If you enter 50, it indicates that 50% of memory is assigned to the Edge transport node.
CPU Reservation (MHz)	Enter CPU reservation in MHz. The maximum amount of MHz is equal to the number of vCPUs multiplied by the normal CPU operation rate of the physical CPU core. **Note** If the MHz value entered exceeds the maximum CPU capacity of the physical CPU cores, the NSX Edge VM might fail to start even though the allocation was accepted.

6 Click **Save**.

If changes made to the resource reservations do not take effect, you might need to reboot the NSX Edge VM from VMware vCenter.

The NSX Edge VM appliance autostarts on ESXi host reboot provided the NSX Edge cluster has vSphere HA turned off. For more details on vSphere HA, see the vSphere documentation.

Replacing NSX Edge Hardware or Redeploying NSX Edge Nodes VM

Know the scenarios to replace an existing NSX Edge node with a physical server or with an NSX Edge VM.

If you encounter any of the following scenarios, you might want to consider replacing or redeploying NSX Edge nodes deployed in your network:

- A Bare Metal server (physical server) used as an NSX Edge node crashed.
- An NSX Edge VM appliance reached defunct status.
- An NSX Edge VM appliance placement (network/datastore/compute, form factor, latency sensitivity, CPU and memory) needs to be changed.

Supported replacement paths are:

Action	Supported Path (From... To)	Reasons
Replace Hardware	From Physical Server or NSX Edge VM to Physical Server	■ Hardware crash ■ New physical server requirement
Redeploy NSX Edge	From Physical Server or NSX Edge VM to NSX Edge VM	■ Change required to compute, storage, network, CPU, memory, latency sensitivity settings of the existing NSX Edge node. ■ Defunct NSX Edge node

Replace NSX Edge Hardware

The need to replace an NSX Edge node arises when an existing physical server fails or when you want to replace the VM form factor to a physical server form factor. But, there could be other reasons for replacing the hardware, such as node reaching defunct status and so on.

Prerequisites

- (Physical servers and NSX Edge VMs manually deployed through vSphere Client) Before you replace a NSX Edge node, ensure connectivity between the existing NSX Edge node and NSX Manager is down. If connectivity is Up, NSX does not allow the existing NSX Edge node to be replaced with a new one.
- Starting with NSX 4.0.1.1, NSX Edge VM hardware version will no longer default to `virtualHW.version` **13**. NSX Edge VM hardware will depend on the underlying version of the ESXi host. VM hardware versions compatible with ESXi hosts are listed in KB article 2007240.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Go to **System → Fabric → Nodes → Edge Transport Node**.

3 Connect to the NSX Manager console using an SSH session.

4 On the NSX Manager appliance, run the `get certificate api thumbprint` command.

The command output is a string of alphanumeric numbers that is unique to this NSX Manager.

For example:

```
NSX-Manager1> get certificate api thumbprint
659442c1435350edbbc0e87ed5a6980d892b9118f851c17a13ec76a8b985f57
```

5 To replace an existing NSX Edge node with a new NSX Edge node, connect to the new NSX Edge node and run the `join management-plane` command.

Provide the following information:

- Hostname or IP address of the NSX Manager with an optional port number
- User name of the NSX Manager

- Certificate thumbprint of the NSX Manager
- Password of the NSX Manager

```
join management-plane <Manager-IP> thumbprint <Manager-thumbprint>
username admin replace
```

Note

- If the old node is an NSX Edge VM, which is deployed through NSX Manager API, then on executing the `join management-plane` cli command, NSX Manager deletes the existing NSX Edge VM and joins the new physical hardware to the management plane. NSX Manager prepares the new physical server as a transport node using the same configuration as that of the old NSX Edge VM.
- If the old node is a physical server, ensure it is powered down before you run the `join management-plane` CLI command. NSX Manager removes the existing physical server and joins the new physical hardware to the management plane. If you power on the old node, run `del nsx` before reusing it your environment.

6 Go to the NSX Manager user interface to view the configuration status of the newly replaced NSX Edge node.

What to do next

If you want to bring up the replaced physical server or manually deployed NSX Edge VM appliance as part of your network, ensure that the node is disconnected from the network. Then, run `del nsx` to completely delete NSX VIBs on the node. See the *NSX Installation Guide* for more details on `del nsx`.

After you run `del nsx` on the host, old entries of logical routers, VTEP IP addresses, uplink IP addresses are released. You can now prepare the replaced physical server as an NSX transport node.

Redeploy an NSX Edge VM Appliance

You want to redeploy an NSX Edge VM when it becomes defunct or its placement in the datacenter needs to change. For example, when the NSX Edge must be moved to another datastore or compute resource, redeploy the NSX Edge node. You can also move the node to another network. However, there could be other reasons to redeploy depending on your network requirements.

You can only redeploy an existing NSX Edge node (physical server or NSX Edge VM appliance) with an NSX Edge VM appliance.

Prerequisites

- While you can change some configurations of NSX Edge transport node payload, do not change these configurations on the existing NSX Edge node, that is to be redeployed by a new node:
  - Failure domain
  - Transport node connectivity
  - Physical NIC configuration
  - Logical routers
  - Load balancer allocations
- Ensure connectivity between NSX Edge node and NSX Manager is down if the existing NSX Edgenode is a physical server or a manually deployed VM through vSphere Client. If connectivity is Up, then NSX does not allow the existing NSX Edge node to be replaced with a new one.
- Existing autodeployed NSX Edge will remain with hardware version 13. Starting with NSX 4.0.1.1, if the NSX Edge VM is redeployed, the new NSX Edge VM is automatically deployed with an upgraded hardware version compatible with the ESXi host version. VM hardware versions compatible with ESXi hosts are listed in KB article 2007240.

Procedure

1. (Physical server or NSX Edge deployed through vSphere Client) Open an SSH session and connect to the NSX Edge console.
2. Verify the logical routes configured on the NSX Edge node through the CLI console, `get logical-routers`.
3. Power off NSX Edge node.
4. Verify that the NSX Edge node is disconnected from NSX Manager by running the following API command.

   `GET api/v1/transport-nodes/<edgenode>/state`

   ```
   "node_deployment_state":
          {"state": MPA_Disconnected"}
   ```

   The `node_deployment_state` value is `MPA Disconnected`, which indicates you can proceed to redeploy the NSX Edge node.

   Note If `node_deployment_state` is `Node Ready`, NSX Manager displays an error `78006 - Manager connectivity to Edge node must be down. Else replacement/ redeployment of hardware is not allowed.`

5 Alternatively, view the state of connectivity between NSX Edge node and NSX Manager from the **Edge Transport Node** page. A disconnected NSX Edge node displays the following system message, `Configuration Error, Edge VM MPA Connectivity is down`.

6 If the NSX Edge node is an auto-deployed node, run `GET /<NSX-Manager-IPaddress>/api/v1/transport-nodes/<edgenode>`. Copy the output payload of this API.


```
"resource_type": "EdgeNode",
            "id": "9f34c0ea-4aac-4b7f-a02c-62f306f96649",
            "display_name": "Edge_TN2",
            "description": "EN",
            "external_id": "9f34c0ea-4aac-4b7f-a02c-62f306f96649",
            "ip_addresses": [
                "10.170.94.240"
            ],
            "_create_user": "admin",
            "_create_time": 1600106319056,
            "_last_modified_user": "admin",
            "_last_modified_time": 1600106907312,
            "_system_owned": false,
            "_protection": "NOT_PROTECTED",
            "_revision": 2
        },
        "is_overridden": false,
        "failure_domain_id": "4fc1e3b0-1cd4-4339-86c8-f76baddbaafb",
        "resource_type": "TransportNode",
        "id": "9f34c0ea-4aac-4b7f-a02c-62f306f96649",
        "display_name": "Edge_TN2",
        "_create_user": "admin",
        "_create_time": 1600106319399,
        "_last_modified_user": "admin",
        "_last_modified_time": 1600106907401,
        "_system_owned": false,
        "_protection": "NOT_PROTECTED",
        "_revision": 1
    }
```

7 You can choose from one of the redeploy scenarios:

Choice	Actions
Redeploy an existing NSX Edge node (physical server or manually deployed node) with an NSX Edge VM node (deployed through NSX Manager API)	Perform the following in the API command, `/api/v1/transport-nodes/<transport-node-id>?action=redeploy` ■ Paste the payload in the body of the redeploy API. ■ Verify the `deployment_config` section references the compute manager, datstore, and network details, where you want to redeploy the node. Ensure these values are consistent with the values used in the `node_settings` section. ■ Add login passwords in the `deployment_config` section. NSX Manager redeploys the NSX Edge node based on the details provided in the `deployment_config` section.
Change the placement of the existing NSX Edge node	Perform the following in the API command, `/api/v1/transport-nodes/<transport-node-id>?action=redeploy` ■ Paste the payload in the body of the redeploy API. ■ In the `deployment_config` section, reference the new compute manager, datastore, network, CPU, memory, or latency sensitivity details.

An example of the `POST https://<manager-ip>/api/v1/transport-nodes/<transport-node-id>?action=redeploy`

```
{
    "node_id": "9f34c0ea-4aac-4b7f-a02c-62f306f96649",
    "host_switch_spec": {
        "host_switches": [
            {
                "host_switch_name": "nsxvswitch_overlay",
                "host_switch_id": "c0a4a83e-c8b8-4324-a4d7-dbbc07b30b53",
                "host_switch_type": "NVDS",
                "host_switch_mode": "STANDARD",
                "host_switch_profile_ids": [
                    {
                        "key": "UplinkHostSwitchProfile",
                        "value": "f9a2a2fa-b49d-498f-abaf-2fdc81917716"
                    },
                    {
                        "key": "LldpHostSwitchProfile",
                        "value": "9e0b4d2d-d155-4b4b-8947-fbfe5b79f7cb"
                    }
                ],
                "pnics": [
                    {
                        "device_name": "fp-eth0",
                        "uplink_name": "uplink1"
                    }
                ],
                "is_migrate_pnics": false,
                "ip_assignment_spec": {
                    "ip_pool_id": "647d9b0d-0143-4903-91f5-930d9ab011e8",
                    "resource_type": "StaticIpPoolSpec"
```

```
                    },
                    "cpu_config": [],
                    "transport_zone_endpoints": [
                        {
                            "transport_zone_id": "0b33b078-6438-4d9b-a1ec-33211fd36822",
                            "transport_zone_profile_ids": [
                                {
                                    "resource_type": "BfdHealthMonitoringProfile",
                                    "profile_id": "52035bb3-ab02-4a08-9884-18631312e50a"
                                }
                            ]
                        },
                        {
                            "transport_zone_id": "a0133574-48de-4e3a-9407-7db1a68bae41",
                            "transport_zone_profile_ids": [
                                {
                                    "resource_type": "BfdHealthMonitoringProfile",
                                    "profile_id": "52035bb3-ab02-4a08-9884-18631312e50a"
                                }
                            ]
                        }
                    ],
                    "vmk_install_migration": [],
                    "pnics_uninstall_migration": [],
                    "vmk_uninstall_migration": [],
                    "not_ready": false
                }
            ],
            "resource_type": "StandardHostSwitchSpec"
        },
        "transport_zone_endpoints": [],
        "maintenance_mode": "DISABLED",
        "node_deployment_info": {
            "deployment_type": "VIRTUAL_MACHINE",
            "deployment_config": {
                "vm_deployment_config": {
                    "vc_id": "cc82da39-b119-4869-a7fe-a54621cb4d3d",
                    "compute_id": "domain-c9",
                    "storage_id": "datastore-14",
                    "host_id": "host-12",
                    "compute_folder_id": "group-v5",
                    "management_network_id": "network-16",
                    "hostname": "EdgeSmallFactor",
                    "data_network_ids": [
                        "5638c577-e142-4a50-aed3-a7079dc3b08c",
                        "5638c577-e142-4a50-aed3-a7079dc3b08c",
                        "5638c577-e142-4a50-aed3-a7079dc3b08c"
                    ],
                    "search_domains": [
                        "eng.vmware.com",
                        "vmware.com"
                    ],
                    "enable_ssh": true,
                    "allow_ssh_root_login": true,
                    "reservation_info": {
```

```
                        "memory_reservation": {
                            "reservation_percentage": 100
                        },
                        "cpu_reservation": {
                            "reservation_in_shares": "HIGH_PRIORITY",
                            "reservation_in_mhz": 0
                        }
                    },
                    "resource_allocation": {
                        "cpu_count": 4,
                        "memory_allocation_in_mb": 8192
                    },
                    "placement_type": "VsphereDeploymentConfig"
                },
                "form_factor": "MEDIUM",
                "node_user_settings": {
                     "cli_username": "admin",
                        "root_password":"Admin!23Admin",
                        "cli_password":"Admin!23Admin"
                }
            },
            "node_settings": {
                "hostname": "EdgeSmallFactor",
                "search_domains": [
                    "eng.vmware.com",
                    "vmware.com"
                ],
                "enable_ssh": true,
                "allow_ssh_root_login": true
            },
            "resource_type": "EdgeNode",
            "id": "9f34c0ea-4aac-4b7f-a02c-62f306f96649",
            "display_name": "Edge_TN2",
            "description": "EN",
            "external_id": "9f34c0ea-4aac-4b7f-a02c-62f306f96649",
            "ip_addresses": [
                "10.170.94.240"
            ],
            "_create_user": "admin",
            "_create_time": 1600106319056,
            "_last_modified_user": "admin",
            "_last_modified_time": 1600106907312,
            "_system_owned": false,
            "_protection": "NOT_PROTECTED",
            "_revision": 2
        },
        "is_overridden": false,
        "failure_domain_id": "4fc1e3b0-1cd4-4339-86c8-f76baddbaafb",
        "resource_type": "TransportNode",
        "id": "9f34c0ea-4aac-4b7f-a02c-62f306f96649",
        "display_name": "Edge_TN2",
        "_create_user": "admin",
        "_create_time": 1600106319399,
        "_last_modified_user": "admin",
        "_last_modified_time": 1600106907401,
```

```
        "_system_owned": false,
        "_protection": "NOT_PROTECTED",
        "_revision": 1
    }
```

Note If the old node is an NSX Edge VM node deployed through NSX Manager UI, then you do not need to provide login credentials in the `node_user_settings` section in the API payload.

See the *NSX API Guide* for more information on the payload details.

8 In NSX Manager, verify the **Configuration Status** of the new NSX Edge node.

9 Alternatively, verify the status of the newly prepared NSX Edge transport node by running the API command, `Get api/v1/transport-nodes/<node-id>/state`.

10 Verify logical router configurations are migrated to the new NSX Edge node, by running the `get logical-routers` CLI command.

11 Verify the TEP address remain same on the replaced NSX Edge node.

12 Verify NSX Edge cluster status is Up. API is `GET api/v1/edge-clusters/<cluster>`. If NSX is configured to use NSX Federation, verify the intersite status us Up.

13 Check NSX Edge transport node and cluster state API to verify status is Up.

14 Troubleshoot error messages:

- `(78006) NSX Manager connectivity to Edge node must be down. Else replacement of hardware is not allowed`: Ensure NSX Edge node is not connected to NSX Manager.
- `(16064) Deployment configuration is missing`: In the redeploy API, enter details for the `deployment_config` section.
- `(16066) Login password is missing`: Provide login credentials.
- `(15019) Insufficient resources on node to be allocated to load balancer pool`: The form factor size of the new NSX Edge node might be smaller than the form factor of the old NSX Edge node. The new form factor might not have enough resources to be allocated to load balancer pool.

What to do next

- If you want to bring up a replaced physical server or manually deployed NSX Edge VM appliance as part of your network, ensure that the node is disconnected from the network. Then, run `del nsx` to completely delete NSX VIBs on the node. See the *NSX Installation Guide* for more details on `del nsx`.

After you run del nsx on the host, old entries of logical routers, VTEP IP addresses, uplink IP addresses are released. You can now prepare the replaced physical server as a new NSX transport node.

- After you redeploy a NSX Edge VM Appliance, few of the security parameters are set to their default values. Reconfigure these parameters as per your environment.
  - `set auth-policy minimum-password-length <password-length-arg>`

    Set the minimum number of characters that passwords must have. The smallest value that can be set is 8

    For example, `nsx> set auth-policy minimum-password-length 12`
  - `set user <node-username> password-expiration <password-expiration-arg>`

    Set number of days the user's password is valid after a password change.

    Where, *<username>* is the Username of user,

    <password-expiration> is the number of days password valid after change (1 - 9999)

    For example, `nsx> set user audit password-expiration 120`
  - `set auth-policy cli max-auth-failures <auth-failures-arg>`

    Set the number of failed CLI authentication attempts that are allowed before the account is locked. If set to 0, account lockout is disabled.

    Where, *<auth-failures>* is the number of authentication failures to trigger lockout

    For example, `nsx> set auth-policy cli max-auth-failures 5`
  - `set banner`

    Set the security banner or message of the day.

    For example, `nsx> set banner`

    `Enter TEXT message. End with 'Ctrl-D'`
  - `reset dataplane hugepage`

    Reset the hugepage-related boot time option to factory default.

    For example

    `nsx-edge-1> reset dataplane hugepage`

    ```
    0000:0b:00.0 already bound to driver vfio-pci, skipping
    0000:1b:00.0 already bound to driver vfio-pci, skipping
    0000:13:00.0 already bound to driver vfio-pci, skipping
    INFO: Config was written to: /config/vmware/edge/config.json
    Generating grub configuration file ...
    Found linux image: /vmlinuz-3.14.17-nn4-server
    Found initrd image: //initrd.img-3.14.17-nn4-server
    ```

```
File descriptor 4 (/tmp/ffinvYglp (deleted)) leaked on lvs invocation. Parent PID
32203: /bin/sh
done
INFO: Updated grub. Please reboot to take effect.
```

Configure Advanced NSX Edge Parameters to Improve Performance

To improve NSX Edge node performance, you can configure the Coalescing Scheme and Coalescing Param advanced configuration parameters, which are available when you prepare an NSX Edge transport node. Call the NSX Edge transport node API to configure the Coalescing configuration.

Coalescing Scheme and Coalescing Param configuration define the virtual network interrupt rate in interrupts per second. Coalescing can reduce the number of interrupts, thus potentially decreasing CPU utilization. Though this could increase network latency, many workloads are not impacted by additional network latency of anywhere from a few hundred microseconds to a few milliseconds, and the reduction in virtual networking overhead can potentially allow more virtual machines on an ESXi host. For more information, see vSphere documentation.

Supported Coalescing Schemes: Decide the type of Coalescing Scheme you want to configure and the Coalescing Parameter value you want to set to configure virtual network interrupts on the Edge node.

- `rbc`: Is rate-based coalescing scheme. It supports values in the range of 100 to 100000 interrupts per second. The default value vSphere takes is 4000 interrupts per second.
- `static`: It sets the number of packets before interrupting the CPU.
- `adapt`: Is an adaptive scheme where vSphere decides the interrupt behavior depending upon the load. If the load is lower, the number of interrupts are higher and if the load is higher, the number of interrupts are lower.
- `disabled`: Is used to disable coalescing scheme and param on the NSX Edge node.

To know more about the coalescing scheme and parameters, see Virtual Network Interrupt Coalescing.

There are a couple of ways to configure Coalescing Scheme and Param configuration.

Prerequisites

- Ensure Coalescing Scheme and Coalescing Param configuration is enabled in vSphere. See vSphere documentation.
- VMXNET3 driver is running on VMs.

Procedure

1 To configure Coalescing Scheme and Param, you can configure the NSX Edge node in a couple of ways:

- (API only) Deploy the NSX Edge node by making an API call.
- Deploy the NSX Edge node from the NSX Manager and call the Redeploy transport node API to update the Coalescing Scheme and Param configuration as part of the NSX Edge node settings.

2 Deploy the NSX Edge node transport node API to deploy the node with Coalescing Scheme and Param.

```
POST https://<nsxManagerIp>/api/v1/transport-nodes/
```

```
{
  "host_switch_spec":{
  "host_switches":[
      ....
      ....
    ],
   "resource_type":"StandardHostSwitchSpec"
   },
   "maintenance_mode":"DISABLED",
   "node_deployment_info":{
    ....
    ....
  },
 "node_settings":{
 "hostname":"edgetwo.com",
 "enable_ssh":true,
 "allow_ssh_root_login":true,
 "enable_upt_mode":false,
 "advanced_configuration":[
 {"key": "advanced-config:coalescingScheme",
  "value": "rbc"
       },
       {
   "key": "advanced-config:coalescingParams",
    "value": "4000"
     }
    ]
   },
 "resource_type":"EdgeNode",
 "display_name":"edge-two",
  "tags":[

 ],
 "_revision":0
 },
 "is_overridden":false,
 "failure_domain_id":"4fc1e3b0-1cd4-4339-86c8-f76baddbaafb",
```

```
 "resource_type":"TransportNode",
 "id":"30b425c5-85ca-4402-9705-88b077d08a06",
 "display_name":"edge-nsx-2"
}
```

3 If you already deployed NSX Edge node from NSX Manager, call the redeploy API to edit the settings.

`POST https://<nsxManagerIp>/api/v1/transport-nodes/<edge-id>?action=redeploy`

```
{
  "transport_node":{
 "resource_type": "TransportNode",
 "display_name": "edge-node5",
 "host_switch_spec": {
 "resource_type": "StandardHostSwitchSpec",
  "host_switches": [
    {
     "host_switch_profile_ids": [
       {
        "value": "e331116d-f59e-4004-8cfd-c577aefe563a",
        "key": "UplinkHostSwitchProfile"
        }
        ],
      "host_switch_name": "nsxvswitch",
      "pnics": [
       {
       "device_name": "fp-eth0",
       "uplink_name": "uplink1"
       }
       ]
       },
      "transport_zone_endpoints": [
       {
        "transport_zone_id": "e14c6b8a-9edd-489f-b624-f9ef12afbd8f"
          }
        ],
        "node_id": "8538f119-ba45-4fb1-9cf1-ee849e4cf168",
        "node_deployment_info": {
        "resource_type": "EdgeNode",
        "id": "8538f119-ba45-4fb1-9cf1-ee849e4cf168",
        "display_name": "edge-node5",

     "node_settings":{
             "hostname":"edgetwo.com",
             "enable_ssh":true,
             "allow_ssh_root_login":true,
             "enable_upt_mode":false,
             "advanced_configuration":[
                {
                   "key": "advanced-config:coalescingScheme",
                    "value": "rbc"
                },
                {
```

```
                    "key": "advanced-config:coalescingParams",
                     "value": "4000"
                }
            ]
          },


        "ip_addresses": [
          "192.168.110.37",
          "192.168.110.38"
        ]
      }
}
```

The Coalescing configuration will fail in the following cases:

- Coalescing scheme is not set and coalescing params is set, then coalescing params will be ignored.
- Coalescing scheme is set and coalescing params is not set then NSX will validate and fail this configuration.

Edit NSX Edge VM Restart Policy

You can set the restart policy for an NSX Edge VM deployed from the vSphere Client at a cluster level using the `VM Override` property.

Restart priority ensure that NSX Edge VM is restarted first and virtual network infrastructure is up and running before the worload sends North-South traffic to the NSX Edge VM.

Note When you deploy an NSX Edge node from NSX Manager, redeploy an NSX Edge or upgrade an NSX Edge, restart priority of NSX Edge VM is automatically set to Highest.

To validate NSX Edge restart policy is set to Highest, go to the

Prerequisites

- NSX Edge node must be deployed. See the *NSX Installation Guide* for more details.

Procedure

1. With admin privileges, log in to the vSphere Client.
2. At the cluster lever, vSphere HA is enabled.
3. To verify the NSX Edge restart policy, perform the following steps:
   a. Select a cluster.
   b. Select the **Configure** tab → **vSphere Availability**.

c Verify that **vSphere HA** is turned on.

d In the **Edit Cluster Settings** window, expand **Host Failure Response** section and verify that **Failure Response** is set to **Restart VMs** and **Default VM Restrart Priority** is set to **Highest**.

Results

When an NSX Edge node fails, it is restarted as per the configured restart priority. If you set the priority to Highest, NSX Edge node comes back up before other hosts in the cluster.

Adding and Removing an ESXi Host Transport Node to and from vCenter Servers

You can move an ESXi host transport node from one VMware vCenter (VC) to another, and also from one NSX Manager cluster to another.

Scenario 1: VC1 connected to NSX Manager cluster 1, and VC2 connected to NSX Manager cluster 2

Assuming ESX1, an ESXi host transport node, is in VC1, you can move it to VC2 by performing the following steps:

1 Uninstall NSX from ESX1.

2 Move ESX1 to VC2.

3 Apply a transport node profile to ESX1.

Scenario 2: Both VC1 and VC2 connected to NSX Manager cluster

Assuming ESX1, an ESXi host transport node, is in VC1, you can move it to VC2 by performing the following steps:

1 Uninstall NSX from ESX1.

2 Move ESX1 to VC2.

3 Apply a transport node profile to ESX1.

Scenario 3: VC1 connected to NSX Manager cluster 1

Assuming ESX1, an ESXi host transport node, is in VC1, you can move it to NSX Manager cluster 2 as a standalone host by performing the following steps:

1 Uninstall NSX from ESX1.

2 Add ESX1 to NSX Manager cluster 2.

Note Starting with NSX 4.0.1.1, the user interface displays a warning, if at any point, your host transport node goes out of sync with the Management Plane. To initiate the resync operation, select **Actions > Sync Transport Node** for your host transport node.

Changing the Distributed Router Interfaces' MAC Address

All logical router interfaces in NSX and NSX-V setups have the same MAC address (02:50:56:56:44:52). Starting with NSX 3.0.2, you can change this address in NSX to avoid issues when migrating VMs from an NSX-V setup to an NSX setup.

Changing the MAC address involves making two API calls.

If you have not created any transport node, make the following GET API call. For example:

```
GET https://10.40.79.126/api/v1/global-configs/RoutingGlobalConfig

Response:
{
  "l3_forwarding_mode" : "IPV4_ONLY",
  "logical_uplink_mtu" : 1500,
  "vdr_mac" : "02:50:56:56:44:77",
  "vdr_mac_nested" : "02:50:56:56:44:52",
  "allow_changing_vdr_mac_in_use" : true,
  "resource_type" : "RoutingGlobalConfig",
  "id" : "49b261fe-f4e4-46ad-958c-da9cb4271e32",
  "display_name" : "49b261fe-f4e4-46ad-958c-da9cb4271e32",
  "_create_user" : "system",
  "_create_time" : 1595313890595,
  "_last_modified_user" : "admin",
  "_last_modified_time" : 1595465694142,
  "_system_owned" : false,
  "_protection" : "NOT_PROTECTED",
  "_revision" : 14
}
```

Take the response of the call, change the `vdr_mac` value, and use it to make the following PUT API call. For example:

```
PUT  https://10.40.79.126/api/v1/global-configs/RoutingGlobalConfig
{
  "l3_forwarding_mode" : "IPV4_ONLY",
  "logical_uplink_mtu" : 1500,
  "vdr_mac" : "02:50:56:56:44:99",
  "vdr_mac_nested" : "02:50:56:56:44:53",
  "allow_changing_vdr_mac_in_use" : true,
  "resource_type" : "RoutingGlobalConfig",
  "id" : "49b261fe-f4e4-46ad-958c-da9cb4271e32",
  "display_name" : "49b261fe-f4e4-46ad-958c-da9cb4271e32",
  "_create_user" : "system",
  "_create_time" : 1595313890595,
  "_last_modified_user" : "admin",
  "_last_modified_time" : 1595465694142,
  "_system_owned" : false,
  "_protection" : "NOT_PROTECTED",
  "_revision" : 14
}

Response:
```

```
{
  "l3_forwarding_mode" : "IPV4_ONLY",
  "logical_uplink_mtu" : 1500,
  "vdr_mac" : "02:50:56:56:44:99",
  "vdr_mac_nested" : "02:50:56:56:44:53",
  "allow_changing_vdr_mac_in_use" : true,
  "resource_type" : "RoutingGlobalConfig",
  "id" : "49b261fe-f4e4-46ad-958c-da9cb4271e32",
  "display_name" : "49b261fe-f4e4-46ad-958c-da9cb4271e32",
  "_create_user" : "system",
  "_create_time" : 1595313890595,
  "_last_modified_user" : "admin",
  "_last_modified_time" : 1595466163148,
  "_system_owned" : false,
  "_protection" : "NOT_PROTECTED",
  "_revision" : 15
}
```


If you have already created transport nodes, make the same GET and PUT API calls, except that for the PUT call, set the parameter `allow_changing_vdr_mac_in_use` to `true`.

Configuring Appliances

Some system configuration tasks must be done using the command line or API.

For complete command line interface information, see the *NSX Command-Line Interface Reference*. For complete API information, see the *NSX API Guide*.

Table 27-2. System configuration commands and API requests.

Task	Command Line (NSX Manager and NSX Edge)	API Request (NSX Manager only)
Set system timezone	`set timezone <timezone>`	`PUT https://<nsx-mgr>/api/v1/node`
Set NTP Server	`set ntp-server <ntp-server>`	`PUT https://<nsx-mgr>/api/v1/node/services/ntp`
Set a DNS server	`set name-servers <dns-server>`	`PUT https://<nsx-mgr>/api/v1/node/network/name-servers`
Set DNS Search Domain	`set search-domains <domain>`	`PUT https://<nsx-mgr>/api/v1/node/network/search-domains`

Note The recommended method to configure an NTP server for all appliances is to configure a node profile. See Configure a Node Profile. If you configure an NTP server individually on an appliance, be sure to configure the same NTP server on all the appliances.

Configuring NTP on Appliances and Transport Nodes

Some features require that NTP be configured on all the components in the NSX environment.

It is highly recommended that you configure NTP for all components, regardless of the features you plan to use.

Note Change NTP time configuration in small increments. Do not change the time in a one large jump. If you change the time in a large increment, you must restart the appliance for correct functionality.

To configure NTP on an appliance, see Configuring Appliances.

To configure NTP for all Manager and Edge nodes, see Configure a Node Profile.

To configure NTP for an ESXi host, see the topic Synchronize ESXi Clocks with a Network Time Server, in VSphere Security.

By default, different NTP services run in different Linux distributions. In all releases, run the `timedatectl` to view the synchronization status. Only one NTP service can be running at a time.

The following Linux distributions are supported with their default NTP services:

Linux Server	NTP Client
RHEL/CentOS 7.6	chronyd.service
RHEL/CentOS 7.7	chronyd.service
RHEL/CentOS 8.0	chronyd.service
RHEL/CentOS 8.2	chronyd.service
Ubuntu 18.04	systemd-timesyncd
Ubuntu 20.04	systemd-timesyncd
SLES12 SP4	systemd-timedated

Add a License Key and Generate a License Usage Report

You can add license keys and generate a license usage report. The usage report is a file in CSV format.

When you install NSX Manager, the default license is NSX for vShield Endpoint. This license never expires but has certain restrictions. For more information about licenses, see License Types. You cannot create or update the following objects:

- Tier-0 and tier-1 logical router
- Tier-0 and tier-1 gateway
- Logical switch
- Layer 2 segment (Note: You can create and update service segments.)
- Distributed firewall
- VPN

- NAT
- Load balancer
- Service Insertion
- NSX Intelligence

If you upgrade from a previous release, both the default vShield Endpoint license and the previous default NSX Data Center Evaluation will be available. Note the following:

- You cannot delete the default vShield Endpoint license key.
- You can delete the previous default NSX Data Center Evaluation license key.
- If you add a new vShield Endpoint license, the default vShield Endpoint license will be hidden. If you remove the new vShield Endpoint license, the default vShield Endpoint license will be available again.
- If you add a new NSX Data Center Evaluation license, the default NSX Data Center Evaluation license, if it exists because of an upgrade, will be permanently deleted. If you remove the new NSX Data Center Evaluation license, you will not have any evaluation license.

Note About the evaluation license: If you install the NSX Data Center Evaluation license, it will be valid for 60 days.

Note the following:

- You have the NSX for vShield Endpoint license and the NSX Data Center Evaluation license only.
  - If the NSX Data Center Evaluation license is valid, it will be used.
  - If the NSX Data Center Evaluation license has expired, the NSX for vShield Endpoint license will be used. (Enforcement will take effect.)
- You have the NSX for vShield Endpoint license, the NSX Data Center Evaluation license, and other licenses.
  - If the NSX Data Center Evaluation license is valid, it and the other licenses will be used.
  - If the NSX Data Center Evaluation license has expired, the other licenses will be used.

If a license has expired or will expire within 60 days, an alarm will be generated each time you log in. You can view alarms by going to **Home > Alarms**.

If you have only the NSX for vShield Endpoint license, after you log in, an informational banner message will let you know that the license has certain restrictions (see above). If you have any expired or expiring license, after you log in, a warning banner message will let you know.

You can add multiple license keys for the same edition of NSX (Formerly known as NSX-T Data Center). It is not necessary to go to https://my.vmware.com to combine the keys.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Select **System > Licenses > Add License**.

3 Enter a license key.

You will get an error message if you enter an invalid license key or enter an add-on license without the required base SKU.

4 To remove a license, select a license and click **Delete**.

You will get an error message if you remove a base SKU and there is an add-on license that depends on it.

5 To generate a license usage report, select **Export > License Usage Report**.

The CSV report lists the VM, CPU, unique concurrent user, vCPU and core usage numbers of the following features:

- Switching and Routing
- NSX Edge load balancer
- VPN
- Distributed Firewall
- Context Aware Micro-Segmentation - Application identification
- Context Aware Micro-Segmentation - Identity firewall for remote desktop session host
- Service Insertion
- Identity Firewall
- Enhanced Guest Introspection
- Micro-Segmentation Planning (L4) (CPU, core, VM, and concurrent user information only)
- NSX Federation (CPU, core, VM, and concurrent user information only)
- Anomaly Detection
- Endpoint Context
- Micro-Segmentation Planning (L7)
- Distributed Load Balancer
- Distributed IDPS
- URL Categorization
- Distributed Malware Prevention

Note The following features are disabled for the Limited Export Release version:

- IPSec VPN
- HTTPS-based Load Balancer

License Types

You must have a base license and then you can add an Add-On license on top of the base license.

Base Licenses

The following base license types are available.

Base Licenses Available Starting with NSX 4.0.1.1:

- NSX Advanced with Threat Prevention
- NSX Advanced with Advanced Threat Prevention
- NSX Enterprise Plus with Threat Prevention
- NSX Enterprise Plus with Advanced Threat Prevention

NSX Data Center:

- NSX Data Center Standard
- NSX Data Center Professional
- NSX Data Center Advanced
- NSX Data Center Enterprise Plus
- NSX Data Center for Remote Office Branch Office (ROBO)
- NSX Data Center Evaluation
- NSX for vSphere - Standard
- NSX for vSphere - Advanced
- NSX for vSphere - Enterprise
- NSX for vShield Endpoint
- VMware Container Networking with Antrea Enterprise (3.2 and later versions)

NSX-T:

The NSX-T licenses are available for 3.0.3.1, 3.1.3.1, 3.2, and later versions.

- NSX-T Standard
- NSX-T Professional
- NSX-T Advanced
- NSX-T Advanced Limited Edition
- NSX-T Enterprise Plus
- NSX-T for Remote Office Branch Office (ROBO)
- NSX-T Evaluation

Limited Export License

The following limited export license types are available:

- VMware NSX Enterprise per Processor (Limited Export)
- NSX Data Center Advanced per Processor (for Limited Export)
- NSX Data Center Evaluation
- NSX for vShield Endpoint
- NSX Data Center Distributed Threat Prevention
- NSX-T Advanced Limited Edition
- NSX-T Evaluation

Note For Limited Export Release version, you can add the NSX Data Center Distributed Threat Prevention add-on license only if the VMware NSX Enterprise per Processor (Limited Export) or NSX Data Center Advanced per Processor (for Limited Export) license exists. You cannot delete the VMware NSX Enterprise per Processor (Limited Export) or NSX Data Center Advanced per Processor (for Limited Export) license until the add-on license is deleted.

Security Licenses

Security licenses that are available in NSX v3.1.0 and later:

- NSX Firewall
- NSX Firewall with Advanced Threat Prevention
- NSX Firewall with Bare Metal
- NSX Firewall Evaluation

NSX Distributed Firewall licenses that are available in NSX v3.2 and later:

- NSX Distributed Firewall
- NSX Distributed Firewall with Threat Prevention
- NSX Distributed Firewall with Advanced Threat Prevention
- NSX Threat Prevention Add-On for Distributed Firewall
- NSX Advanced Threat Prevention Add-On for Distributed Firewall, NSX-T Data Center Advanced or NSX-T Data Center Enterprise Plus

NSX Gateway Firewall licenses that are available in NSX v3.2 and later:

- NSX Gateway Firewall – VM
- NSX Gateway Firewall – ISO
- NSX Gateway Firewall with Threat Prevention – VM
- NSX Gateway Firewall with Threat Prevention – ISO

- NSX Gateway Firewall with Advanced Threat Prevention – VM
- NSX Gateway Firewall with Advanced Threat Prevention – ISO
- NSX Threat Prevention Add-On for Gateway Firewall – VM
- NSX Threat Prevention Add-On for Gateway Firewall – ISO
- NSX Advanced Threat Prevention Add-On for Gateway Firewall – VM
- NSX Advanced Threat Prevention Add-On for Gateway Firewall – ISO

See also the section "Base Licenses Available Starting with NSX 4.0.1.1" at the top of this page where the base license is bundled with either Threat Prevention or Advanced Threat Prevention.

VMware Cloud Provider Program (VCPP) Licenses

VMware Cloud Provider Program (VCPP) licenses that are available in v3.1.1 and later:

- NSX Data Center SP Base
- NSX Data Center SP Professional
- NSX Data Center SP Advanced
- NSX Data Center SP Enterprise Plus

Add-On Licenses

You must apply all add-on license on top of one of the corresponding base license as listed in the following table.

Note You cannot delete the base license until the add-on license is deleted.

Table 27-3. Add-On Licenses (Before v3.2)

Add-On License	Base License
NSX Data Center Distributed Threat Prevention (Distributed IDS included) (Introduced in v3.0)	One of the following base license is required: ■ NSX Data Center Advanced ■ NSX Data Center Enterprise Plus ■ NSX Firewall ■ NSX-T Advanced ■ NSX-T Enterprise Plus
NSX Advanced Threat Prevention (Introduced in v3.1.1)	One of the following base license is required: ■ NSX Data Center Advanced ■ NSX Data Center Enterprise Plus ■ NSX Firewall ■ NSX Advanced ■ NSX Enterprise Plus

Table 27-4. Add-On Licenses (From v3.2 and later)

Add-On License	Base License
NSX Threat Prevention Add-On for NSX Distributed Firewall	One of the following base license is required: ■ NSX Distributed Firewall ■ NSX Data Center Advanced ■ NSX Data Center Enterprise Plus ■ NSX-T Advanced ■ NSX-T Enterprise Plus
NSX Advanced Threat Prevention Add-On for NSX Distributed Firewall, NSX Data Center Advanced, NSX Data Center Enterprise Plus	One of the following base license is required: ■ NSX Data Center Advanced ■ NSX Data Center Enterprise Plus ■ NSX Distributed Firewall ■ NSX Distributed Firewall with Threat Prevention
NSX Threat Prevention Add-On for Gateway Firewall - VM	NSX Gateway Firewall - VM
NSX Advanced Threat Prevention Add-On for Gateway Firewall - VM	One of the following base licenses is required: ■ NSX Gateway Firewall - VM ■ NSX Gateway Firewall with Threat Prevention - VM
NSX-T Advanced Add-On	One of the following base licenses is required: ■ NSX Distributed Firewall and NSX Threat Prevention Add-On for Distributed Firewall (Need both) ■ NSX Distributed Firewall with Advanced Threat Prevention ■ NSX Distributed Firewall and NSX Advanced Threat Prevention Add-On for Distributed Firewall, NSX Data Center Advanced or NSX Data Center Enterprise Plus (Need both)
NSX-T Enterprise Plus Add-On	One of the following base license is required: ■ NSX Advanced Threat Prevention Add-On for Distributed Firewall, NSX Data Center Advanced or NSX Data Center Enterprise Plus ■ NSX Data Center Distributed Threat Prevention ■ NSX Advanced Threat Prevention ■ NSX Threat Prevention Add On for NSX Distributed Firewall

Standards Compliance

NSX can be configured to use FIPS 140-2 validated cryptographic modules to run in FIPS-compliant mode. The modules are validated to FIPS 140-2 standards by the NIST Cryptographic Module Validation Program (CMVP).

All exceptions to FIPS compliance can be retrieved using the compliance report. See View Compliance Status Report for more information.

The following validated modules are used:

- VMware's BoringCrypto Module 3.0: Certificate #4028
- VMware's OpenSSL FIPS Object Module version 2.0.20-vmw: Certificate #3857

- BC-FJA (Bouncy Castle FIPS Java API) version 1.0.2.1: Certificate #3673
- VMware's IKE Crypto Module version 1.1.0: Certificate #3435
- VMware's VPN Crypto Module version 2.0: Certificate #4286

You can find more information about the cryptographic modules that VMware has validated against the FIPS 140-2 standard here: https://www.vmware.com/security/certifications/fips.html.

By default, load balancer uses modules that have FIPS mode turned off. You can turn on FIPS mode for the modules used by load balancer. See Configure Global FIPS Compliance Mode for Load Balancer for more information.

View Compliance Status Report

You can view a compliance report for NSX features. You can use the report to configure your NSX environment to adhere to your IT policies and industry standards.

The compliance report includes information about each non-compliant configuration.

Table 27-5. Compliance Report Information

Compliance Report Column	Description	Example
Non Compliance Code	Code to identify the type of non-compliance.	72301
Description	Description of the type of non-compliance.	Certificate is not CA signed.
Resource Name	Name or ID of the affected resource.	nsx-manager-1
Resource Type	Type of resource affected.	CertificateComplianceReporter
Affected Resources	Number of affected resources. The number can be 0 if there are non-compliant configurations present, but the feature is not used.	1

You can also retrieve the report using the API: `GET /policy/api/v1/compliance/status`.

Procedure

1. From your browser, log in with admin privileges to an NSX Manager at https://<nsx-manager-ip-address>.
2. From the **Home** page, click **Monitoring Dashboards > Compliance Report**.

Compliance Status Report Codes

You can find the information about the meaning of the compliance report codes in the compliance status report.

For a complete list of events, see the NSX Event Catalog.

Table 27-6. Compliance Report Codes

Code	Description	Compliance Status Source	Remediation
72001	Encryption is deactivated.	Reports this status if a VPN IPSec Profile configuration contains `NO_ENCRYPTION`, `NO_ENCRYPTION_AUTH_AES_GMAC_128`, `NO_ENCRYPTION_AUTH_AES_GMAC_192`, or `NO_ENCRYPTION_AUTH_AES_GMAC_256` encryption_algorithms. This status affects IPSec VPN session configurations which use the reported non-compliant configurations.	Add a VPN IPSec Profile that uses compliant encryption algorithms and use the profile in all VPN configurations. See Add IPSec Profiles.
72011	BGP messages with neighbor bypass integrity check. No message authentication defined.	Reports this status if BGP neighbors has no confgured password. This status affects the BGP neighbor configuration.	Configure a password on the BGP neighbor and update the tier-0 gateway configuration to use the password. See Configure BGP.
72012	Communication with BGP neighbor uses weak integrity check. MD5 is used for message authentication.	Reports this status if MD5 authentication is used for the BGP neighbor password. This status affects the BGP neighbor configuration.	No remediation available as NSX supports only MD5 authentication for BGP.
72021	SSL version 3 used for establishing secure socket connection. It is recommended to run TLSv 1.1 or higher and fully disable SSLv3 that have protocol weaknesses.	Reports this status if SSL version 3 configuration is in the load balancer client SSL profile, load balancer server SSL profile, or load balancer HTTPS monitor. This status affects the following configurations: ■ Load balancer pools that are associated with HTTPS monitors. ■ Load balancer virtual servers that are associated with load balancer client SSL profiles or server SSL profiles.	Configure an SSL profile to use TLS 1.1 or later and use this profile in all load balancer configurations. See Add an SSL Profile.

Table 27-6. Compliance Report Codes (continued)

Code	Description	Compliance Status Source	Remediation
72022	TLS version 1.0 used for establishing secure socket connection. Run TLSv 1.1 or higher and fully disable TLSv1.0 that have protocol weaknesses.	Reports this status if the load balancer client SSL profile, load balancer server SSL profile, or load balancer HTTPS monitor includes the TLSv1.0 configuration. This status affects the following configurations: ■ Load balancer pools that are associated with HTTPS monitors. ■ Load balancer virtual servers that are associated with load balancer client SSL profiles or server SSL profiles.	Configure an SSL profile to use TLS 1.1 or later and use this profile in all load balancer configurations. See Add an SSL Profile.
72023	Weak Diffie-Hellman group is used.	Reports this error if a VPN IPSec Profile or VPN IKE Profile configuration includes the following Diffie-Hellman groups: 2, 5, 14, 15 or 16. Groups 2 and 5 are weak Diffie-Hellman groups. Groups 14, 15, and 16 are not weak groups, but are not FIPS-compliant. This status affects IPSec VPN session configurations which use the reported non-compliant configurations.	Configure the VPN Profiles to use Diffie-Hellman group 19, 20, or 21. See Adding Profiles.
72024	Load balancer FIPS global setting is deactivated.	Reports this error if the load balancer FIPS global setting is deactivated. This status affects all load balancer services.	Enable FIPS for load balancer. See Configure Global FIPS Compliance Mode for Load Balancer.
72025	QAT running on Edge node is Non-FIPS Compliant.	Quick Assist Technologies (QAT) is a set of hardware accelerated services provided by Intel for cryptography and compression.	
72026	Bouncy-castle FIPS module is not ready.		Check to ensure your applications are running and check your logs.

Table 27-6. Compliance Report Codes (continued)

Code	Description	Compliance Status Source	Remediation
72200	Insufficient true entropy available.	Reports this status if a pseudo random number generator generates the entropy rather than relying on hardware-generated entropy. NSX Manager node does not use hardware-generated entropy because is does not have the required hardware acceleration support to create sufficient true entropy.	Use newer hardware to run the NSX Manager node. Most recent hardware supports this feature. **Note** If the underlying infrastructure is virtual, you will not get true entropy.
72201	Entropy source unknown.	Reports this status when no entropy status is available for the indicated node.	This error is an internal communication error preventing the NSX Manager from being able to determine the source of entropy. Open a service request with VMware.
72301	Certificate is not CA signed.	Reports this status when one of the NSX Manager certificates is not CA signed. NSX Manager uses the following certificates: ■ Syslog certificate. ■ API certificates for the individual NSX Manager nodes. ■ Cluster certificate used for the NSX ManagerVIP.	Install CA-signed certificates. Refer Chapter 23 Certificates.
72302	Certificate is missing Extended Key Usage information.	Extended Key Usage (EKU) extensions present in the CSR should also be present in the server certificates.	The Extended Key Usage (EKU) needs to match the intended usage. For example, SERVER when connecting to a server, CLIENT when used as a client certificate on a server.
72303	Certificate has been revoked.		Validity of the certificate is in a terminal state and cannot recover.
72304	Certificate is not RSA.		Ensure your certificate public key is for an RSA certificate.

Table 27-6. Compliance Report Codes (continued)

Code	Description	Compliance Status Source	Remediation
72305	Certificate is not Elliptic Curve.	Reports this status when your certificate fails EAL4+ compliance.	Replace your non-compliant certificates with CA-signed certificates. See Replace Certificates .
72306	Certificate is missing Basic Constraints.		The certificate does not appear to be valid for the selected purpose or may be missing necessary fields or values..
72307	Unable to fetch the CRL.		
72308	The CRL is invalid.		Check CRL to determine why it has been marked invalid.
72309	The Corfu Certificate Expiry Check has been disabled.		
72310	Some Compute Manager does not have RSA certificate or the certificate key length is less than 3072 bits.	When a compute manager (for example, vCenter) is connected to the NSX Manager, it passes its certificate to NSX Manager. If the certificate is not an RSA type, or if the certificate is an RSA type, but the certificate's RSA key size is less then 3072 bits, it triggers this error.	Delete the compute manager from NSX Manager. Replace the compute manager's certificate with an RSA certificate with key size of at least 3072 bits.
72401	Gateway TLS Inspection is non-FIPS compliant.		
72402	System not FIPS compliant.		
72403	Principal Identity presence detected in system. Remove all the Principal Identities for EAL4 compliance.		
72404	OIDC end-point presence detected in system. Remove all the OIDC end-points for EAL4 compliance.		

Table 27-6. Compliance Report Codes (continued)

Code	Description	Compliance Status Source	Remediation
72501	Configured user passwords are not compliant. Users must use high-quality passwords of length at least 16 characters.	Reports when user passwords fail EAL4+ compliance.	The password for local users (except for the root user) must be at least 16 characters. In other words, this means that the admin user password must be at least 16 characters. This is in addition to the password complextity check that is required when modifying the password.
72502	Unable to get synchronized users.		
72503	SSH service for Manager appliance found to be in running state.		
72504	Non-admin active user accounts detected.	When any user other than admin is active in the NSX Manager.	Deactivate all other users other than admin.
73000	Error occurred while gathering platform compliance data.		

Configure Global FIPS Compliance Mode for Load Balancer

There is a global setting for FIPS compliance for load balancers. By default, the setting is turned off to improve performance.

Changing the global configuration for FIPS compliance for load balancers affects new load balancer instances, but does not affect any existing load balancer instances.

If the global setting for FIPS for load balancer (`lb_fips_enabled`) is set to *true*, new load balancer instances use modules that comply with FIPS 140-2. Existing load balancer instances might be using non-compliant modules.

To make the change take effect on existing load balancers, you must detach and reattach the load balancer from the tier-1 gateway.

You can check the global FIPS compliance status for load balancer using `GET /policy/api/v1/compliance/status`.

```
...
{
    "non_compliance_code": 72024,
    "description": "Load balancer FIPS global setting is disabled.",
    "reported_by": {
        "target_id": "971ca477-df1a-4108-8187-7918c2f8c3ba",
        "target_display_name": "971ca477-df1a-4108-8187-7918c2f8c3ba",
```

```
                "target_type": "FipsGlobalConfig",
                "is_valid": true
            },
            "affected_resources": [
                {
                    "path": "/infra/lb-services/LB_Service",
                    "target_id": "/infra/lb-services/LB_Service",
                    "target_display_name": "LB_1",
                    "target_type": "LBService",
                    "is_valid": true
                }
            ]
        },
        ...
```

Note The compliance report displays the global setting for FIPS compliance for load balancer. Any given load balancer instance can have a FIPS compliance status that is different from the global setting.

Procedure

1 Retrieve the global FIPS setting for load balancer.

```
GET https://nsx-mgr1/policy/api/v1/infra/global-config
```

Example response body:

```
{
    "fips": {
        "lb_fips_enabled": false
    },
    "resource_type": "GlobalConfig",
    "id": "global-config",
    "display_name": "global-config",
    "path": "/infra/global-config",
    "relative_path": "global-config",
    "marked_for_delete": false,
    "_create_user": "system",
    "_create_time": 1561225479619,
    "_last_modified_user": "admin",
    "_last_modified_time": 1561937915337,
    "_system_owned": true,
    "_protection": "NOT_PROTECTED",
    "_revision": 2
}
```

2 Change the global FIPS setting for load balancer.

The global setting is used when you create new load balancer instances. Changing the setting does not affect existing load balancer instances.

```
PUT https://nsx-mgr1/policy/api/v1/infra/global-config
```

Example request body:

```
{
    "fips": {
        "lb_fips_enabled": true
    },
    "resource_type": "GlobalConfig",
    "_revision": 2
}
```

Example response body:

```
{
    "fips": {
        "lb_fips_enabled": true
    },
    "resource_type": "GlobalConfig",
    "id": "global-config",
    "display_name": "global-config",
    "path": "/infra/global-config",
    "relative_path": "global-config",
    "marked_for_delete": false,
    "_create_user": "system",
    "_create_time": 1561225479619,
    "_last_modified_user": "admin",
    "_last_modified_time": 1561937960950,
    "_system_owned": true,
    "_protection": "NOT_PROTECTED",
    "_revision": 3
}
```

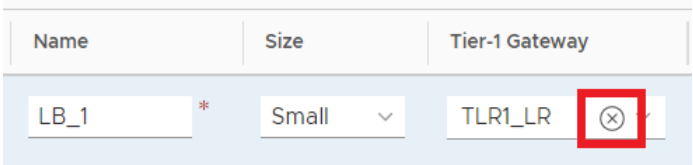
3 If you want any existing load balancer instances to use this global setting, you must detach and reattach the load balancer from the tier-1 gateway.

Caution Detaching a load balancer from the tier-1 gateway results in a traffic interruption for the load balancer instance.

a Navigate to **Networking > Load Balancing**.

b On the load balancer you want to detach, click the three dots menu (⋮), then click **Edit**.

c Click ⊗, then click **Save** to detach the load balancer from the tier-1 gateway.



d Click the three dots menu (), then click **Edit**.

e Select the correct gateway from the **Tier-1 Gateway** drop-down menu, then click **Save** to reattach the load balancer to the tier-1 gateway.

Common Criteria Compliance

NSX is designed to be EAL4+ compliant in accordance with the Common Criteria Certification Program.

For more information about Common Criteria, see the Common Criteria Portal.

If your environment is not compliant with EAL4+, NSX will raise alarms. For more information about the alarms, see the "Security Compliance Events" section in NSX Event Catalog.

NSX supports the following security functional requirements:

Requirement	Description	Auditable
FAU_GEN.1	Audit data generation	Yes
FAU_SAR.1	Audit review	No
FAU_STG.1	Protected audit trail storage	No
FCS_CKM.1/TLS	Cryptographic key generation (for asymmetric keys)	No
FCS_COP.1/TLS.HMAC	Cryptographic operation (hashing)	No
FCS_COP.1/TLS	Cryptographic operation (TLS)	No
FCS_CKM.2/TLS	Cryptographic key distribution (TLS)	Yes
FCS_CKM.4	Cryptographic key destruction	No
FCS_RNG.1/OSSL	Random number generation (OpenSSL)	No
FCS_RNG.1/BC	Random number generation (Bouncy Castle)	No
FDP_IFC.1	Subset information flow control	No
FDP_IFF.1	Simple security attributes	Yes
FIA_AFL.1	Authentication failure handling	Yes
FIA_SOS.1	Verification of secrets	Yes
FIA_UAU.2	User authentication before any action	Yes
FIA_UID.1	Timing of identification	Yes
FMT_SMR.1	Security roles	Yes
FMT_SMF.1	Specification of management functions	Yes
FMT_MOF.1	Management of security functions behavior	No

Requirement	Description	Auditable
FMT_MSA.1	Management of security attributes	No
FMT_MSA.3	Static attribute initialisation	No
FPT_TDC.1	Inter-TSF basic TSF data consistency	Yes
FTP_ITC.1	Inter-TSF trusted channel	Yes

For each requirement, events that are auditable will be logged. For more information about event logging, see Log Messages and Error Codes.

Collect Support Bundles

You can collect support bundles on registered cluster and fabric nodes and download the bundles to your machine or upload them to a file server.

If you choose to download the bundles to your machine, you get a single archive file consisting of a manifest file and support bundles for each node. If you choose to upload the bundles to a file server, the manifest file and the individual bundles are uploaded to the file server separately.

NSX Cloud Note If you want to collect the support bundle for CSM, log in to CSM, go to **System > Utilities > Support Bundle** and click **Download**. The support bundle for PCG is available from NSX Manager using the following instructions. The support bundle for PCG also contains logs for all the workload VMs.

If you want to collect support bundles for Antrea container clusters that are registered to NSX, see *Collect Support Bundles for an Antrea Container Cluster* in the *NSX Administration Guide*.

NSX Application Platform Note For information about collecting support bundles for NSX Application Platform, see the *Deploying and Managing the VMware NSX Application Platform* documentation.

Procedure

1. From your browser, log in as a local admin user to an NSX Manager at https://*nsx-manager-ip-address/login.jsp?local=true*.
2. Select **System > Support Bundle**
3. Select the target nodes.

   The available types of nodes are **Management Nodes**, **Edges**, **Hosts**, and **Public Cloud Gateways**.
4. (Optional) Specify log age in days to exclude logs that are older than the specified number of days.

5 (Optional) Toggle the switch that indicates whether to include or exclude core files and audit logs.

Note Core files and audit logs might contain sensitive information such as passwords or encryption keys.

6 (Optional) Select the check box to upload the bundles to a remote file server.

7 Click **Start Bundle Collection** to start collecting support bundles.

Depending on how many log files exist, each node might take several minutes.

8 Monitor the status of the collection process.

The status tab shows the progress of collecting support bundles.

9 Click **Download** to download the bundle if the option to send the bundle to a file remote server was not set.

The bundle collection may fail for a manager node if there is not enough disk space. If you encounter an error, check whether older support bundles are present on the failed node. Log in to the NSX Manager UI of the failed manager node using its IP address and initiate the bundle collection from that node. When prompted by the NSX Manager, either download the older bundle or delete it.

Understanding Support Bundle File Paths

Know the different files collected when you download a support bundle and the features that store logs in these file paths.

NSX Node Support Bundle

File Path	Features Using the File Path
/var/log/syslog	All
/var/log/nsx-audit.log	All
/var/log/policy/policy.log	All
/var/log/policy/policy_restart.log	All
/var/log/policy/policy-tomcat-wrapper.log	All
/var/log/proton/activity-stats.log	All
/var/log/proton/gc.log	All
/var/log/proton/localhost_access_log.txt	All
/var/log/proton/nsxapi.log	All
/var/log/proton/proton_restart.log	All

File Path	Features Using the File Path
/var/log/proton/proton-tomcat-wrapper.log	All
/var/log/proton-tomcat/nsxapi.log	All
/var/log/proxy/gc.log	AAA (Auth)
/var/log/proxy/keystore.log	AAA (Auth)
/var/log/proxy/localhost.log	AAA (Auth)
/var/log/proxy/localhost_access_log.txt	AAA (Auth)
/var/log/proxy/proxy-tomcat-wrapper.log	AAA (Auth)
/var/log/proxy/reverse-proxy.log	AAA (Auth)
/var/log/phonehome-coordinator/localhost_access_log.txt	Alarm Event Framework, Telemetry, Notification Framework, Capacity Dashboard
/var/log/phonehome-coordinator/phonehome-audit-log.txt.0	Alarm Event Framework, Telemetry, Notification Framework, Capacity Dashboard
/var/log/phonehome-coordinator/phonehome-coordinator.log	Alarm Event Framework, Telemetry, Notification Framework, Capacity Dashboard
/var/log/cbm/cbm.log	Clustering
/var/log/cm-inventory/localhost_access_log.txt	CM-inventory (Compute Manager)
/var/log/vmware/top-cpu.log	Common - Appliance Health
/var/log/vmware/top-mem.log	Common - Appliance Health
/var/log/stats	Common - IP, memory, thread stats
/var/log/cloudnet/nsx-ccp-transaction.log	Controller
/var/log/cloudnet/nsx-ccp.log	Controller
/var/log/core	Cores generated
/var/log/corfu/corfu.9000.log	Corfu DB
/var/log/corfu/corfu-compactor-audit.log	Corfu DB
/var/log/corfu-nonconfig/corfu.9040.log	Corfu DB
/var/log/corfu-nonconfig/nonconfig-corfu-compactor-audit.log	Corfu DB
/var/log/cloudnet/nsx-ccp.log	DNS, Upgrade, IPDiscovery, Groups
/var/log/cm-inventory/cm-inventory.log	Endpoint Security, Deployment, NSX Edge Install, Vmotion, Inventory, NCP/Openshift/PKS-windows, Install/ Host Upgrade, Install
/var/log/migration-coordinator/v2t/cm.log	Endpoint Security (V2T Migration), V2T

File Path	Features Using the File Path
/var/log/async-replicator-service/ar.log	Federation (Local Manager)
/var/log/proton/nsx-common-dashboard.log	NCP/Openshift/PKS-windows
/var/log/search/elasticsearch_index_indexing_slowlog.log	NCP/Openshift/PKS-windows
/var/log/search/search-policy.log	NCP/Openshift/PKS-windows
/var/run/log/nsxaVim.log	Stateless Hosts/ Transport Node Profile
/var/log/policy/localhost_access_log.txt	Policy realization
/var/log/search/elasticsearch.log	UI
/var/log/intelligence-upgrade-coordinator/localhost_access_log.txt	Upgrade
/var/log/upgrade	Upgrade
/var/log/resume-upgrade.log	Upgrade
/var/log/upgrade-coordinator/localhost_access_log.txt	Upgrade
/var/log/upgrade-coordinator/upgrade-coordinator.log	Upgrade
/var/log/nvpapi/api_server.log	Upgrade, BnR
/var/log/upgrade-coordinator/localhost_access_log.txt	Upgrade Coordinator
/var/log/migration-coordinator/v2t/summary.log	V2T (NSX-V to NSX)
/var/log/migration-coordinator/migration-coordinator.log	V2T (NSX-V to NSX)

NSX Edge Support Bundle

File Path	Features Using the File Path
/var/log/syslog	All
/var/log/pcm/pcm.log	CSM
/var/log/rsyslog.log	CSM
/var/log/upgrade/history_date\<>\<>.log	CSM
/controller/data/data_dump	DHCP
/controller/falcon/falcon_dump	DHCP
/controller/span/config_span_graph_dump	DHCP
/controller/span/dag_span_graph_dumpContainer_Dfw_Canary	DHCP
/controller/mediator/mediator_dump	DHCP

File Path	Features Using the File Path
/var/log/cloudnet/nsx-ccp.log	DHCP
/var/log/proton/nsxapi.log	DHCP
/edge/dns_fdr-all	DNS
/edge/dns-forwarder	DNS
/edge/service_binding	DNS
/var/log/join_mp.log	Edge Install, Install, Upgrade
/edge/fw-if-ruleset/	Load Balancer
/edge/Lb_UUID/	Load Balancer
/edge/lb-*	Load Balancer
/var/log/lb/<LB_UUID>/logs/error.log	Load Balancer
/var/log/lb/<LB_UUID>/logs/ acess_vipid.log	Load Balancer
/var/log/lb/<LB_UUID>/logs/nginx.conf	Load Balancer
/var/log/lb/<LB_UUID>/lbconf_gen.log	Load Balancer
/var/log/core/	Load Balancer, L3VPN, L2VPN, Edge HA, Upgrade
/var/log/nsx-cli/nsxcli.log	L3VPN,L2VPN,Edge HA, Upgrade
/var/log/rcpm/frr-config.log	L3VPN, Upgrade
/var/log/rcpm/frr-reload.log	L3VPN, Upgrade
/var/log/vmware/top-cpu.log	NCP/Openshift/PKS-windows
/var/log/vmware/top-mem.log	NCP/Openshift/PKS-windows
/var/log/frr/frr.log	Openstack, Upgrade
/opt/vmware/nsx-nestdb/bin/nestdb-cli	Openstack, Upgrade
/var/log/nginx/access.log	Openstack
/var/log/upgrade	Upgrade
/var/log/nvpapi/api_server.log	Upgrade, BnR

NSX Application Platform Support Bundle

The following support bundles include log information for the features that NSX Application Platform hosts. These features are NSX Intelligence, NSX Network Detection and Response, NSX Malware Prevention, and NSX Metrics.

File Path	Features Using the File Path
/var/log/napp/supportbundle/platform_service_dbg.log	NSX Application Platform
/var/log/napp/supportbundle/ <K8s_worker_node_name>/kafka*.log	NSX Application Platform
/var/log/napp/supportbundle/ <K8s_worker_node_name>/zookeeper*.log	NSX Application Platform
/var/log/napp/supportbundle/ <K8s_worker_node_name>/ spark-app-rawflow-driver*.log	NSX Application Platform
/var/log/napp/supportbundle/ <K8s_worker_node_name>/authserver*.log	NSX Application Platform
/var/log/napp/supportbundle/ <K8s_worker_node_name>/cluster-api*.log	NSX Application Platform
/var/log/napp/supportbundle/ <K8s_worker_node_name>/common-agent*.log	NSX Application Platform
/var/log/napp/supportbundle/ <K8s_worker_node_name>/projectcontour*.log	NSX Application Platform
/var/log/napp/supportbundle/ <K8s_worker_node_name>/routing-controller*.log	NSX Application Platform
/var/log/napp/supportbundle/ <K8s_worker_node_name>/trust-manager*.log	NSX Application Platform
/var/log/napp/supportbundle/ <K8s_worker_node_name>/upgrade-coordinator*.log	NSX Application Platform
/var/log/napp/supportbundle/ <K8s_worker_node_name>/postgresql*.log	NSX Application Platform
/var/log/napp/supportbundle/ <K8s_worker_node_name>/visualization*.log	NSX Intelligence Visualization
/var/log/napp/supportbundle/ <K8s_worker_node_name>/recommendation*.log	NSX Intelligence Recommendation
/var/log/napp/supportbundle/ <K8s_worker_node_name>/rawflowcorrelator*.log	NSX Intelligence
/var/log/napp/supportbundle/ <K8s_worker_node_name>/ overflowcorrelator*.log	NSX Intelligence
/var/log/napp/supportbundle/ <K8s_worker_node_name>/*druid*.log	NSX Intelligence
/var/log/napp/supportbundle/ <K8s_worker_node_name>/nsx-config*.log	NSX Intelligence
/var/log/napp/supportbundle/ <K8s_worker_node_name>/redis*.log	NSX Intelligence
/var/log/napp/supportbundle/ <K8s_worker_node_name>/minio*.log	NSX Intelligence
/var/log/napp/supportbundle/ <K8s_worker_node_name>/nsx-ndr*.log	NSX Network Detection and Response

File Path	Features Using the File Path
/var/log/napp/supportbundle/<K8s_worker_node_name>/sa-*.log	NSX Malware Prevention
/var/log/napp/supportbundle/<K8s_worker_node_name>/malware*.log	NSX Malware Prevention
/var/log/napp/supportbundle/<K8s_worker_node_name>/metrics*.log	NSX Metrics

NSX Global Manager Support Bundle

File Path	Features Using the File Path
/var/logdata-log.log	Federation
/var/log/async-replicator- ar.log	Federation
/var/log/connections.log	Federation
/var/log/gm.log	Federation
/var/log/nvpapi/api_access.log	Federation
/var/log/global-manager/gmanager.log	Grouping

Debugging NSX at Runtime

The Online Diagnostic System (ODS) feature automates debugging of NSX at runtime. ODS is implemented through runbooks that come in-built in NSX. Runbooks contain debugging procedures and have full observability of NSX components. A runbook generates a debugging report. It also generates runtime artifacts of an issue such as packet capture, live core dump of user process, and output of scripts and tools. These artifacts can be collected later and used for offline analysis and debugging. Note that you cannot modify a predefined runbook.

NSX comes with the following predefined runbooks that can be run only on ESXi:

- PNIC Performance runbook - This runbook can identify pNIC TX and RX performance issues.
- Overlay Tunnel runbook – This runbook can identify overlay tunnel failures, such as gateway configuration error tunnel missing, tunnel down.
- Port Block runbook – This runbook can identify causes for a port's blockage, such as a DVPort could be blocked because of incorrect LSP/LS configuration.
- ADF Collection runbook - This runbook gathers ESXi datapath performance data through ADF collector.
- Controller Connectivity runbook – This runbook can identify controller connectivity issues caused by controller failure, proxy failure, underlay network outage, or FQDN resolution failure.

You can use APIs to perform following runbook operations:

- Invoke a runbook to initiate the runtime debugging
- Check the debugging report
- Download artifacts generated by the invoked runbook
- Manage the lifecycle of a runbook

To debug at runtime, perform the following steps:

Step 1: Fetch a list of predefined runbook.

Run the following API to fetch a list of predefined runbooks.

```
GET https://<nsx-mgr>/policy/api/v1/infra/sha/pre-defined-runbooks
```

This API returns a list of predefined runbooks along with the following information:

- Configuration details.
- Node type on which the runbook is supported.
- General details of the runbook such as id, name, and path.
- Parameter details if any required at the time of invoking the runbook.

Example Response:


```
{
    "results": [
        {
            "version": {
                "major": 1,
                "minor": 0
            },
            "default_config": {
                "enabled": true,
                "timeout": 300,
                "threshold_number": 5,
                "throttle_cycle": 10
            },
            "supported_node_types": [
                "nsx-esx"
            ],
            "parameters": [
                {
                    "name": "advanced",
                    "optional": true,
                    "parameter_type": "BOOLEAN",
                    "default_value": "False"
                },
                {
                    "name": "cycle",
                    "optional": false,
                    "parameter_type": "INTEGER",
                    "max": "20",
```

```
                    "min": "1"
                },
                {
                    "name": "interval",
                    "optional": true,
                    "parameter_type": "INTEGER",
                    "max": "300",
                    "min": "1"
                }
            ],
            "resource_type": "OdsPredefinedRunbook",
            "id": "00000000-0000-4164-6643-6f6c6c656374",
            "display_name": "AdfCollect",
            "path": "/infra/sha/pre-defined-runbooks/00000000-0000-4164-6643-6f6c6c656374",
            "relative_path": "00000000-0000-4164-6643-6f6c6c656374",
            "parent_path": "/infra",
            "remote_path": "",
            "unique_id": "069a4c7b-532a-4926-a402-6a7986a306b2",
            "realization_id": "069a4c7b-532a-4926-a402-6a7986a306b2",
            "owner_id": "f2a2a9e1-2578-435d-877d-47e01eb04954",
            "origin_site_id": "f2a2a9e1-2578-435d-877d-47e01eb04954",
            "marked_for_delete": false,
            "overridden": false,
            "_create_time": 1669600880454,
            "_create_user": "system",
            "_last_modified_time": 1669600880454,
            "_last_modified_user": "system",
            "_system_owned": false,
            "_revision": 0
        },
        {
            "version": {
                "major": 1,
                "minor": 0
            },
            "default_config": {
                "enabled": true,
                "timeout": 60,
                "threshold_number": 5,
                "throttle_cycle": 3
            },
            "supported_node_types": [
                "nsx-esx"
            ],
            "resource_type": "OdsPredefinedRunbook",
            "id": "0000436f-6e74-726f-6c6c-6572436f6e6e",
            "display_name": "ControllerConn",
            "path": "/infra/sha/pre-defined-runbooks/0000436f-6e74-726f-6c6c-6572436f6e6e",
            "relative_path": "0000436f-6e74-726f-6c6c-6572436f6e6e",
            "parent_path": "/infra",
            "remote_path": "",
            "unique_id": "3914cbe4-41b4-45f1-9dad-2bd96a2de0d8",
            "realization_id": "3914cbe4-41b4-45f1-9dad-2bd96a2de0d8",
            "owner_id": "f2a2a9e1-2578-435d-877d-47e01eb04954",
            "origin_site_id": "f2a2a9e1-2578-435d-877d-47e01eb04954",
```

```
            "marked_for_delete": false,
            "overridden": false,
            "_create_time": 1669600880493,
            "_create_user": "system",
            "_last_modified_time": 1669600880493,
            "_last_modified_user": "system",
            "_system_owned": false,
            "_revision": 0
        },
        {
            "version": {
                "major": 1,
                "minor": 0
            },
            "default_config": {
                "enabled": true,
                "timeout": 120,
                "threshold_number": 5,
                "throttle_cycle": 5
            },
            "supported_node_types": [
                "nsx-esx"
            ],
            "parameters": [
                {
                    "name": "src",
                    "optional": false,
                    "parameter_type": "COMPOUND"
                },
                {
                    "name": "dst",
                    "optional": false,
                    "parameter_type": "COMPOUND"
                }
            ],
            "resource_type": "OdsPredefinedRunbook",
            "id": "0000004f-7665-726c-6179-54756e6e656c",
            "display_name": "OverlayTunnel",
            "path": "/infra/sha/pre-defined-runbooks/0000004f-7665-726c-6179-54756e6e656c",
            "relative_path": "0000004f-7665-726c-6179-54756e6e656c",
            "parent_path": "/infra",
            "remote_path": "",
            "unique_id": "3597af28-e670-456e-8347-4d1a53a5cb90",
            "realization_id": "3597af28-e670-456e-8347-4d1a53a5cb90",
            "owner_id": "f2a2a9e1-2578-435d-877d-47e01eb04954",
            "origin_site_id": "f2a2a9e1-2578-435d-877d-47e01eb04954",
            "marked_for_delete": false,
            "overridden": false,
            "_create_time": 1669600880518,
            "_create_user": "system",
            "_last_modified_time": 1669600880518,
            "_last_modified_user": "system",
            "_system_owned": false,
            "_revision": 0
        },
```

```
        {
            "version": {
                "major": 1,
                "minor": 0
            },
            "default_config": {
                "enabled": true,
                "timeout": 120,
                "threshold_number": 5,
                "throttle_cycle": 5
            },
            "supported_node_types": [
                "nsx-esx"
            ],
            "parameters": [
                {
                    "name": "lsp",
                    "optional": false,
                    "parameter_type": "STRING"
                }
            ],
            "resource_type": "OdsPredefinedRunbook",
            "id": "00000000-0000-0000-506e-696350657266",
            "display_name": "PnicPerf",
            "path": "/infra/sha/pre-defined-runbooks/00000000-0000-0000-506e-696350657266",
            "relative_path": "00000000-0000-0000-506e-696350657266",
            "parent_path": "/infra",
            "remote_path": "",
            "unique_id": "53f29b77-dcf5-4561-85ec-8f35280e3f3a",
            "realization_id": "53f29b77-dcf5-4561-85ec-8f35280e3f3a",
            "owner_id": "f2a2a9e1-2578-435d-877d-47e01eb04954",
            "origin_site_id": "f2a2a9e1-2578-435d-877d-47e01eb04954",
            "marked_for_delete": false,
            "overridden": false,
            "_create_time": 1669600880553,
            "_create_user": "system",
            "_last_modified_time": 1669600880553,
            "_last_modified_user": "system",
            "_system_owned": false,
            "_revision": 0
        },
        {
            "version": {
                "major": 1,
                "minor": 0
            },
            "default_config": {
                "enabled": true,
                "timeout": 60,
                "threshold_number": 5,
                "throttle_cycle": 3
            },
            "supported_node_types": [
                "nsx-esx"
            ],
```

```
            "parameters": [
                {
                    "name": "vif",
                    "optional": false,
                    "parameter_type": "STRING"
                }
            ],
            "resource_type": "OdsPredefinedRunbook",
            "id": "00000000-0000-0050-6f72-74426c6f636b",
            "display_name": "PortBlock",
            "path": "/infra/sha/pre-defined-runbooks/00000000-0000-0050-6f72-74426c6f636b",
            "relative_path": "00000000-0000-0050-6f72-74426c6f636b",
            "parent_path": "/infra",
            "remote_path": "",
            "unique_id": "6f411a92-30c0-4838-9758-c00220cb5fab",
            "realization_id": "6f411a92-30c0-4838-9758-c00220cb5fab",
            "owner_id": "f2a2a9e1-2578-435d-877d-47e01eb04954",
            "origin_site_id": "f2a2a9e1-2578-435d-877d-47e01eb04954",
            "marked_for_delete": false,
            "overridden": false,
            "_create_time": 1669600880572,
            "_create_user": "system",
            "_last_modified_time": 1669600880572,
            "_last_modified_user": "system",
            "_system_owned": false,
            "_revision": 0
        }
    ],
    "result_count": 5,
    "sort_by": "display_name",
    "sort_ascending": true
}
```


In the above response, the following predefined runbooks are returned.

- AdfCollect
- ControllerConn
- OverlayTunnel
- PnicPerf
- PortBlock

  If a runbook requires a parameter at the time of invocation, the parameter details are specified in the *parameters* key. For example, the Overlay Tunnel runbook requires two parameters, `source` and `destination` that are local and remote VTEP IPs of the tunnel to be diagnosed.

Step 2: Invoke the runbook.

Run the following API to invoke a runbook.

```
POST https://<nsx-mgr>/policy/api/v1/infra/sha/runbook-invocations/<invoke-name>
```

Example Request: Invoking the Overlay Tunnel runbook.

```
POST https://<nsx-mgr>/policy/api/v1/infra/sha/runbook-invocations/OverlayTunnel
{
    "runbook_path": "/infra/sha/pre-defined-runbooks/0000004f-7665-726c-6179-54756e6e656c",
    "target_node": "6c7a9374-459d-46b2-9ea6-c63b37c7cc38",
    "arguments": [
        {
            "key": "src",
            "value": "192.168.0.11"
        },
        {
            "key": "dst",
            "value": "192.168.0.10"
        }
    ]
}
```

Step 3: Check the report of the invoked runbook

Run the following API to get the report of the invoked runbook.

```
GET https://<nsx-mgr>/policy/api/v1/infra/sha/runbook-invocations/<invoke-name>/
report
```

Example Request:

```
GET https://<nsx-mgr>/policy/api/v1/infra/sha/runbook-invocations/OverlayTunnel/
report
```

Example Request:

```
GET https://<nsx-mgr>/policy/api/v1/infra/sha/runbook-invocations/OverlayTunnel/
report
```

Example Response:

```
{
    "invocation_id": "70527fed-1e5e-4fed-a880-28cde04a66b1",
    "target_node": "6c7a9374-459d-46b2-9ea6-c63b37c7cc38",
    "timestamp": 1662469099,
    "sys_info": {
        "host_name": "sc2-10-185-4-158.eng.vmware.com",
        "os_name": "VMkernel",
        "os_version": "7.0.3"
    },
    "result_message": "Tunnel 192.168.0.11 -> 192.168.0.10 is in up state",
    "recommendation_message": "No changes required as the tunnel is in up state.",
    "step_details": [
        {
            "step_id": 1,
            "action_summary": "Check the status of tunnel for the given source/destination
VTEPs.",
            "action_result": "Tunnel 192.168.0.11 -> 192.168.0.10 is in state up"
        }
```

```
    ],
    "status": {
        "request_status": "SUCCESS",
        "operation_state": "FINISHED"
    }
}
```

The response returns some metadata information, such as timestamp and system details like host name and operating system. The report also returns the result of debugging, remediation suggestion if any along with steps executed for debugging with action summary and action result of each step. If debugging is interrupted for any reason, the `operation_state` field will hold a value to define the reason for the interruption. In case the runbook invocation does not succeed, the report provides error details and does not show debugging related fields.

Step 4: Download artifacts.

1 Run the following API to gather artifacts.

```
POST https://{{mgr-ip}}/policy/api/v1/infra/sha/runbook-invocations/{{invocation-name}}
```

Example Request:

```
POST https://{{mgr-ip}}/policy/api/v1/infra/sha/runbook-invocations/{{invocation-name}}
{
    "runbook_path": "/infra/sha/pre-defined-runbooks/
00000000-0000-4164-6643-6f6c6c656374",
    "target_node": "{{target-node}}",
    "arguments": [
        {
            "key": "advanced",
            "value": "{{advanced-mode}}"
        },
        {
            "key": "cycle",
            "value": "{{cycle-count}}"
        },
        {
            "key": "interval",
            "value": "{{interval-in-sec}}"
        }
    ]
}
```

If the `advanced` parameter is set to false at the time of invocation, the runbook will collect topology info, net-stats, NSX DP stats and uplink info. If the `advanced` parameter is set to true, the runbook will additionally run advanced performance tools such as vmkstats (available only on the physical machine).

The `cycle` parameter defines the number of times the ADF collector is executed in an invocation.

The `interval` parameter defines the waiting interval between consecutive ADF collector executions. It takes effect only when the `cycle` parameter is set to greater than 1.

Example Response :


```
{
    "invocation_id": "80a0037a-52e1-48d7-b28e-c3bfb8475e8c",
    "target_node": "b794f78f-7eb0-433f-8f11-63e6b3121c28",
    "timestamp": 1668674073,
    "sys_info": {
        "host_name": "sc2-rdops-vm06-dhcp-204-101.eng.vmware.com",
        "os_name": "VMkernel",
        "os_version": "7.0.3"
    },
    "result_message": "ADF data collection runbook completes.",
    "recommendation_message": "No action needs to be taken.",
    "step_details": [
        {
            "step_id": 1,
            "action_summary": "Run ADF collector.",
            "action_result": "ADF data collection is successfully performed in the
following time point(s) along with the following artifact(s): [(2022-11-17 08:35:30,
a44f0446-0ac3-4e7f-8513-fb1248985d9e.tar)]",
            "artifacts": [
                "a44f0446-0ac3-4e7f-8513-fb1248985d9e"
            ]
        }
    ],
    "status": {
        "request_status": "SUCCESS",
        "operation_state": "FINISHED"
    }
}
```


2 Run the following API to download the artifacts.

`GET https://<nsx-mgr>/policy/api/v1/infra/sha/runbook-invocations/<invoke-name>/artifact`

If the runbook generated the artifacts, the API returns a bundled file, else it returns an error message. Save the binary response to a tar.gz file. This file contains a runbook invocation report (in JSON) as well as a tar file for the performed ADF collection.

Changing a runbook configuration

A transport node (TN) group is first created and then a runbook profile is bound to it. If you change the configuration of a runbook, it is changed at all nodes to which the runbook profile is bound. You can configure whether a runbook is enabled or not, the debugging timeout, and frequency at which the runbook can be invoked using the throttle cycle mechanism. The throttle cycle mechanism specifies number of times a runbook can be executed in a specific time.

Note that a runbook can have only one profile, but a node might have multiple runbook profiles based on the TN groups to which it belongs. In this case, the profile with the highest priority is applied on the node.

To change the a runbook configuration, perform the following steps:

Step 1: Create a TN group of ESXi.

Run the following API to create a TN group.

```
PATCH https://<nsx-mgr>/policy/api/v1/infra/domains/default/groups/<group-name>
```

Example Request:

```
PATCH https://<nsx-mgr>/policy/api/v1/infra/domains/default/groups/<group-name>
{
    "expression": [
        {
            "paths": [
                "/infra/sites/default/enforcement-points/default/host-transport-nodes/TN1"
            ],
            "resource_type": "PathExpression"
        }
    ],
    "extended_expression": [],
    "reference": false,
    "group_type": [],
    "resource_type": "Group"
}
```

Step 2: Change the configuration of the runbook profile and bind the TN group with the profile.

Run the following API to change the configuration,

```
PATCH https://<nsx-mgr>/policy/api/v1/infra/sha/runbook-profiles/<profile-name>
```

Example Request:

```
PATCH https://<nsx-mgr>/policy/api/v1/infra/sha/runbook-profiles/<profile-name>
{
    "runbook_path": "/infra/sha/pre-defined-runbooks/0000004f-7665-726c-6179-54756e6e656c",
    "applied_to_group_path": "/infra/domains/default/groups/tngroup2",
    "config": {
        "enabled": true,
        "timeout": 120,
        "threshold_number": 2,
        "throttle_cycle": 6
    }
}
```

In this example, the `throttle_cycle` is 6 minutes and the `threshold_number` is 2, which means that within 6 minutes, the runbook can be invoked no more than two times.

For complete information about ODS APIs, see *NSX Intelligence & NSX Application Platform API Guide*.

Log Messages and Error Codes

NSX components write to log files in the directory `/var/log`. On NSX appliances, NSX syslog messages conform with RFC 5424. On ESXi hosts, syslog messages conform with RFC 3164.

Viewing Logs

On NSX appliances syslog messages are in `/var/log/syslog`.

On NSX appliances, you can run the following NSX CLI command to view the logs:

```
get log-file <auth.log | controller | controller-error | http.log | kern.log | manager.log |
node-mgmt.log | policy.log | syslog> [follow]
```

The log files are:

Name	Description
auth.log	Authorization log
controller	Controller log
controller-error	Controller error log
http.log	HTTP service log
kern.log	Kernel log
manager.log	Manager service log
node-mgmt.log	Node management log
nsx-audit-write.log	NSX audit write log
nsx-audit.log	NSX audit log
policy.log	Policy service log
syslog	System log

On hypervisors, you can use Linux commands such as `tac`, `tail`, `grep`, and `more` to view the logs.

Each syslog message has the component (`comp`) and sub-component (`subcomp`) information to help identify the source of the message.

NSX produces logs with facility `local6`, which has a numerical value of 22.

The audit log is part of syslog. An audit log message can be identified by the string `audit="true"` in the `structured-data` field. You can configure an external log server to receive log messages. You can also access audit logs using the API `/api/v1/administration/audit-logs`. The file `nsx-audit.log` contains syslog messages with `audit="true"` in the `structured-data` field. The file `nsx-audit-write.log` contains syslog messages with both `audit="true"` and `update="true"` in the `structured-data` field.

Each syslog and audit log message contains a timestamp generated by an NTP server, if configured, or by the system clock. An example of an audit log message:

```
<182>1 2020-05-05T00:29:02.900Z nsx-manager1 NSX 147324 - [nsx@6876 audit="true"
comp="nsx-manager" level="INFO" reqId="fe75651d-c3e7-4680-8753-9ae9d92d7f0c" subcomp="policy"
username="admin"] UserName="admin", ModuleName="AAA", Operation="GetCurrentUserInfo",
Operation status="success"
```

An example of a syslog message:

```
2023-01-25T21:58:42.173Z w4-mydevice-220.eng.vmware.com NSX 1463241016 - [nsx@6876 comp="nsx-
controller" level="INFO" subcomp="corfu-cluster"] Received maintenance mode status: {da3a0242-
bf6d-ee98-224c-b9799cb5e29f=MAINTENANCE_MODE_OFF}
```

An API call can come from NSX Manager, a policy API client, or an NSX node. All API calls are subject to authentication and authorization, and will generate audit logs. This logging is enabled by default and cannot be disabled. An audit log that is associated with an API call has the following information:

- An entity ID parameter `entId` to identify the object of the API.
- A request ID parameter `req-id` to identify a specific API call.
- An external request ID parameter `ereqId` if the API call contains the header `X-NSX-EREQID:<string>`.
- An external user parameter `euser` if the API call contains the header `X-NSX-EUSER:<string>`.

An audit log message from a policy or manager API call will have the following additional fields. Note that node API (NAPI) and CLI audit logs will not have these fields.

- An `update` flag that shows whether the API operation is a read (GET) or write (PUT/POST/DELETE/...) operation.
- An `operation name` field that shows the name of the API operation.
- An `operation status` field that shows whether the API operation succeeded or failed.
- A `new value` field that shows all parameter values of the API request.

NSX does not have the concept of a privileged mode. API calls from all sources and users are audited.

An example of login and logout syslog messages showing a successful login, a failed login, and logins from 2 different devices (note the different IP addresses):

```
2020-07-07T16:33:20.339Z svc.nsxmanager NSX 1513 SYSTEM [nsx@6876 audit="true"
comp="nsx-manager" level="INFO" subcomp="http"] UserName="admin@10.166.61.56",
ModuleName="ACCESS_CONTROL", Operation="LOGIN", Operation status="success"

2020-07-07T16:33:58.779Z svc.nsxmanager NSX 1513 SYSTEM [nsx@6876 audit="true" comp="nsx-
manager" level="INFO" subcomp="http"] UserName="admin", ModuleName="ACCESS_CONTROL",
Operation="LOGOUT", Operation status="success"
```

```
2020-07-07T16:50:21.301Z svc.nsxmanager NSX 1513 SYSTEM [nsx@6876 audit="true"
comp="nsx-manager" level="INFO" subcomp="http"] UserName="admin@10.166.61.80",
ModuleName="ACCESS_CONTROL", Operation="LOGIN", Operation status="success"

2020-07-07T16:43:20.339Z svc.nsxmanager NSX 1513 SYSTEM [nsx@6876 audit="true"
comp="nsx-manager" level="INFO" subcomp="http"] UserName="admin@10.166.61.56",
ModuleName="ACCESS_CONTROL", Operation="LOGIN", Operation status="failure"
```

An example of a syslog message of a policy API call:

```
<182>1 2020-07-06T18:09:14.210Z svc.nsxmanager NSX 2326 FABRIC [nsx@6876 audit="true"
comp="nsx-manager" entId="68d5a9d0-4691-4c9c-94ed-64fd1c96150f" level="INFO" reqId="4c2335aa-
c973-4f74-983f-331a4f7041ca" subcomp="manager" update="true" username="admin"]
UserName="admin", ModuleName="TransportZone", Operation="CreateTransportZone", Operation
status="success", New
value=[{"transport_type":"OVERLAY","host_switch_name":"nsxvswitch","host_switch_mode":"STANDAR
D","nested_nsx":false,"is_default":false,"display_name":"1-
transportzone-1307","_protection":"UNKNOWN"}]
```

An example of syslog messages of CLI access:

```
2020-07-07T16:36:41.783Z svc.nsxmanager NSX 21018 - [nsx@6876 comp="nsx-manager"
subcomp="cli" username="admin" level="INFO"] NSX CLI started (Manager, Policy, Controller)
for user: admin
2020-07-07T16:36:53.469Z svc.nsxmanager NSX 21018 - [nsx@6876 comp="nsx-manager"
subcomp="cli" username="admin" level="INFO"] NSX CLI stopped for user: admin
```

An example of a syslog message when a user runs a CLI command (in this example, `set user admin password-expiration 100`):

```
<182>1 2020-07-22T20:51:49.017Z manager2 NSX 1864 - [nsx@6876 comp="nsx-manager"
subcomp="cli" username="admin" level="INFO" audit="true"] CMD: set user admin password-
expiration 100 (duration: 2.185s), Operation status: CMD_EXECUTED
```

An example of a syslog message of an NAPI call:

```
<182>1 2020-07-21T21:01:38.803Z manager2 NSX 4690 - [nsx@6876 comp="nsx-manager"
subcomp="node-mgmt" username="admin" level="INFO" audit="true"] admin 'GET /api/v1/node/
services/syslog/exporters' 200 731 "" "PostmanRuntime/7.26.1" 0.004588
```

An example of a syslog message of a CLI command:

```
<182>1 2020-07-21T20:54:40.018Z manager2 NSX 16915 - [nsx@6876 comp="nsx-manager"
subcomp="cli" username="admin" level="INFO" audit="true"] CMD: set logging-server 1.1.1.1
proto udp level info (duration: 4.356s), Operation status: CMD_EXECUTED
```

RFC 5424 and RFC 3164 define the following severity levels:

Severity Level	Description
0	Emergency: system is unusable
1	Alert: action must be taken immediately
2	Critical: critical conditions
3	Error: error conditions
4	Warning: warning conditions
5	Notice: normal but significant condition
6	Informational: informational messages
7	Debug: debug-level messages

All logs with a severity of emergency, alert, critical, or error contain a unique error code in the structured data portion of the log message. The error code consists of a string and a decimal number. The string represents a specific module.

Failure to Access a Log File or Remote Log Server

If NSX fails to access or write messages to a log file, an alarm will be generated. The possible errors are:

- A local log file is missing.
- A local log file's permission or ownership setting prevents NSX from writing to the file.
- NSX is unable to send log messages to a third-party remote log server. Note that an alarm will not be raised if NSX fails to send logs to the Aria Operations for Logs agent.

The alarm can be resolved through the alarm framework.

Log Message Formats

For more information about RFC 5424, see https://tools.ietf.org/html/rfc5424. For more information about RFC 3164, see https://tools.ietf.org/html/rfc3164.

RFC 5424 defines the following format for log messages:

```
<facility * 8 + severity> version UTC-TZ hostname APP-NAME procid MSGID [structured-data] msg
```

A sample log message:

```
<187>1 2016-03-15T22:53:00.114Z nsx-manager NSX - SYSTEM [nsx@6876 comp="nsx-manager"
errorCode="MP4039" subcomp="manager"] Connection verification failed for broker
'10.160.108.196'. Marking broker unhealthy.
```

Error Codes

For a list of error codes, see the knowledge base article 71077 NSX-T Data Center 2.x Error Codes.

Configure Remote Logging

You can configure NSX appliances and hypervisors to send log messages to a remote log server.

Remote logging is supported on NSX Manager, NSX Edge, and hypervisors. You must configure remote logging on each node individually.

For the protocol parameter, the options are UDP, TCP, LI, and the secure protocols TLS and LI-TLS. Note that LI and LI-TLS are Aria Operations for Logs protocols. An Aria Operations for Logs log server supports all the protocols. The protocols LI and LI-TLS can only be used if the log server is Aria Operations for Logs. The benefit of using LI or LI-TLS is that they optimize network usage. If the log server is Aria Operations for Logs, using LI or LI-TLS is recommended. If LI cannot be used, TCP has the advantage of being more reliable, whereas UDP has the advantage of requiring less system and network overhead.

Prerequisites

- Familiarize yourself with the CLI command `set logging-server`. For more information, see the *NSX Command-Line Interface Reference*.
- If you specify the secure protocol TLS or LI-TLS, the server and client certificates must be stored in `/image/vmware/nsx/file-store` on each NSX appliance. Note that certificates in the file store are needed only if the exporter is configured using NSX CLI. If you use the API, then there is no need to use the file store. Once you complete the syslog exporter configuration, you must delete all certificates and keys from this location to avoid potential security vulnerabilities.
- To configure a secure connection to a log server, verify that the server is configured with a CA-signed certificates. For example, if you have a Aria Operations for Logs server `vrli.prome.local` as the log server, you can run the following command from a client to see the certificate chain on the server:

```
root@caserver:~# echo -n | openssl s_client -connect vrli.prome.local:443  | sed -ne '/
^Certificate chain/,/^---/p'
depth=2 C = US, L = California, O = GS, CN = Orange Root Certification Authority
verify error:num=19:self signed certificate in certificate chain
Certificate chain
 0 s:/C=US/ST=California/L=HTG/O=GSS/CN=vrli.prome.local
   i:/C=US/L=California/O=GS/CN=Green Intermediate Certification Authority
 1 s:/C=US/L=California/O=GS/CN=Green Intermediate Certification Authority
   i:/C=US/L=California/O=GS/CN=Orange Root Certification Authority
 2 s:/C=US/L=California/O=GS/CN=Orange Root Certification Authority
   i:/C=US/L=California/O=GS/CN=Orange Root Certification Authority
---
DONE
```

Procedure

1 To configure remote logging on an NSX appliance:

a Run the following command to configure a log server and the types of messages to send to the log server. Multiple facilities or message IDs can be specified as a comma delimited list, without spaces.

```
set logging-server <hostname-or-ip-address[:port]> proto <proto> level <level>
[facility <facility>] [messageid <messageid>] [serverca <filename>] [clientca
<filename>] [certificate <filename>] [key <filename>] [structured-data <structured-
data>]
```

You can run the command multiple times to add multiple configurations. For example:

```
nsx> set logging-server 192.168.110.60 proto udp level info facility local6 messageid
SYSTEM,FABRIC
nsx> set logging-server 192.168.110.60 proto udp level info facility auth,user
```

To forward only audit logs to the remote server, specify `audit="true"` in the `structured-data` parameter. For example:

```
set logging-server <server-ip> proto udp level info structured-data audit="true"
```

All NSX logs use facility `local6`. You should use the `messageid` and `structured-data` filters only when the facility filter is not set or when `local6` is included in the specified facilities.

b Configure secure remote login:

- To configure secure remote logging using the protocol LI-TLS, specify the parameter `proto li-tls`. For example:

```
set logging-server vrli.prome.local proto li-tls level info
messageid SWITCHING,ROUTING,FABRIC,SYSTEM,POLICY,HEALTHCHECK,SHA,MONITORING
serverca intermed-ca-full-chain.crt
```

If the configuration is successful, you will get a prompt without any text. To see the content of the server certificate chain (intermediate followed by root), log in as `root` and run the following command:

```
root@nsx1:~# keytool -printcert -file /image/vmware/nsx/file-store/intermed-ca-
full-chain.crt
Certificate[1]:
Owner: CN=Green Intermediate Certification Authority, O=GS, L=California, C=US
Issuer: CN=Orange Root Certification Authority, O=GS, L=California, C=US
Serial number: 3e726e7fbb3b0a7a6b4edd767f867fd2
Valid from: Sun Mar 15 00:00:00 UTC 2020 until: Mon Mar 17 00:00:00 UTC 2025
Certificate fingerprints:
  MD5:  94:C8:9F:92:56:60:EB:DB:ED:4B:11:17:33:27:C0:C9
  SHA1: 42:9C:3C:51:E8:8E:AC:2E:5E:62:95:82:D7:22:E0:FB:08:B8:64:29
  SHA256:
```

```
58:B8:63:3D:0C:34:35:39:FC:3D:1E:BA:AA:E3:CE:A9:C0:F3:58:53:1F:AD:89:A5:01:0D:D3:89
:9E:7B:C5:69
Signature algorithm name: SHA256WITHRSA
Subject Public Key Algorithm: 4096-bit RSA key
Version: 3
Certificate[2]:
Owner: CN=Orange Root Certification Authority, O=GS, L=California, C=US
Issuer: CN=Orange Root Certification Authority, O=GS, L=California, C=US
Serial number: 3e726e7fbb3b0a7a6b4edd767f867fd1
Valid from: Mon Mar 16 07:16:07 UTC 2020 until: Fri Mar 10 07:16:07 UTC 2045
Certificate fingerprints:
  MD5:  ED:AC:F1:7F:88:05:83:2A:83:C0:09:03:D5:00:CA:7B
  SHA1: DC:B5:3F:37:DF:BD:E0:5C:A4:B7:F4:4C:96:12:75:7A:16:C7:61:37
  SHA256:
F2:5B:DE:8A:F2:31:9D:E6:EF:35:F1:30:6F:DA:05:FF:92:B4:15:96:AA:82:67:E3:3C:C1:69:A3
:E5:27:B9:A5
Signature algorithm name: SHA256WITHRSA
Subject Public Key Algorithm: 4096-bit RSA key
Version: 3
```

The logs for both successful and failure conditions are in `/var/log/loginsight-agent/liagent_2020-MM-DD-<file-num>.log`. If the configuration is successful, you can view the Aria Operations for Logs configuration with the following command:

```
root@nsx1:/image/vmware/nsx/file-store# cat /var/lib/loginsight-agent/liagent-
effective.ini
; Dynamic file representing the effective configuration of VMware Log Insight
Agent (merged server-side and client-side configuration)
;     DO NOT EDIT THIS FILE BY HAND -- YOUR CHANGES WILL BE OVERWRITTEN
; Creation time: 2020-03-22T19:41:21.648800

[server]
hostname=vrli.prome.local
proto=cfapi
ssl=yes
ssl_ca_path=/config/vmware/nsx-node-api/syslog/bb466082-996f-4d77-
b6e3-1fa93f4a20d4_ca.pem
ssl_accept_any_trusted=yes
port=9543
filter={filelog; nsx-syslog; pri_severity <= 6 and ( msgid == "SWITCHING"
or msgid == "ROUTING" or msgid == "FABRIC" or msgid == "SYSTEM" or msgid
== "POLICY" or msgid == "HEALTHCHECK" or msgid == "SHA" or msgid == "MONITORING" )}

[filelog|nsx-syslog]
directory=/var/log
include=syslog;syslog.*
parser=nsx-syslog_parser

[parser|nsx-syslog_parser]
base_parser=syslog
extract_sd=yes

[update]
auto_update=no
```

- To configure secure remote logging using the protocol TLS, specify the parameter `proto tls`. For example:

```
set logging-server vrli.prome.local proto tls level info serverca Orange-
CA.crt.pem clientca Orange-CA.crt.pem certificate gc-nsxt-mgr-full.crt.pem key gc-
nsxt-mgr.key.pem
```

  Note the following:

  - For the `serverCA` parameter, only the root certificate is required, not the full chain.
  - If `clientCA` is different from `serverCA`, only the root certificate is required.
  - The certificate should hold the full chain of the NSX Manager (they should be NDcPP compliant - EKU, BASIC and CDP (CDP - this check can be ignored)).

  You can inspect the content of each certificate. For example:

```
root@gc3:~# keytool -printcert -file /image/vmware/nsx/file-store/Orange-CA.crt.pem
Owner: CN=Orange Root Certification Authority, O=GS, L=California, C=US
Issuer: CN=Orange Root Certification Authority, O=GS, L=California, C=US
Serial number: 3e726e7fbb3b0a7a6b4edd767f867fd1
Valid from: Mon Mar 16 07:16:07 UTC 2020 until: Fri Mar 10 07:16:07 UTC 2045
Certificate fingerprints:
  MD5:  ED:AC:F1:7F:88:05:83:2A:83:C0:09:03:D5:00:CA:7B
  SHA1: DC:B5:3F:37:DF:BD:E0:5C:A4:B7:F4:4C:96:12:75:7A:16:C7:61:37
  SHA256:
F2:5B:DE:8A:F2:31:9D:E6:EF:35:F1:30:6F:DA:05:FF:92:B4:15:96:AA:82:67:E3:3C:C1:69:A3
:E5:27:B9:A5
Signature algorithm name: SHA256WITHRSA
Subject Public Key Algorithm: 4096-bit RSA key
Version: 3
root@gc3:~#

root@gc3:/image/vmware/nsx/file-store# keytool -printcert -file gc-nsxt-mgr-
full.crt.pem
Certificate[1]:
Owner: CN=gc.prome.local, O=GS, L=HTG, ST=California, C=US
Issuer: CN=Green Intermediate Certification Authority, O=GS, L=California, C=US
Serial number: bdf43ab31340b87f323b438a2895a075
Valid from: Mon Mar 16 07:26:51 UTC 2020 until: Wed Mar 16 07:26:51 UTC 2022
Certificate fingerprints:
        MD5:  36:3C:1F:57:96:07:84:C0:6D:B7:33:9A:8D:25:4D:27
        SHA1: D1:4E:F9:45:2D:0D:34:79:D2:B4:FA:65:28:E0:5C:DC:74:50:CA:3B
        SHA256:
3C:FF:A9:5D:AA:68:44:44:DD:07:2F:DD:E2:BE:9C:32:19:7A:03:D5:26:8D:5F:AD:56:CA:D2:6C
:91:96:27:6F
Signature algorithm name: SHA256WITHRSA
Subject Public Key Algorithm: 4096-bit RSA key
Version: 3
Certificate[2]:
Owner: CN=Green Intermediate Certification Authority, O=GS, L=California, C=US
Issuer: CN=Orange Root Certification Authority, O=GS, L=California, C=US
Serial number: 3e726e7fbb3b0a7a6b4edd767f867fd2
```

```
Valid from: Sun Mar 15 00:00:00 UTC 2020 until: Mon Mar 17 00:00:00 UTC 2025
Certificate fingerprints:
        MD5:  94:C8:9F:92:56:60:EB:DB:ED:4B:11:17:33:27:C0:C9
        SHA1: 42:9C:3C:51:E8:8E:AC:2E:5E:62:95:82:D7:22:E0:FB:08:B8:64:29
        SHA256:
58:B8:63:3D:0C:34:35:39:FC:3D:1E:BA:AA:E3:CE:A9:C0:F3:58:53:1F:AD:89:A5:01:0D:D3:89
:9E:7B:C5:69
Signature algorithm name: SHA256WITHRSA
Subject Public Key Algorithm: 4096-bit RSA key
Version: 3
Certificate[3]:
Owner: CN=Orange Root Certification Authority, O=GS, L=California, C=US
Issuer: CN=Orange Root Certification Authority, O=GS, L=California, C=US
Serial number: 3e726e7fbb3b0a7a6b4edd767f867fd1
Valid from: Mon Mar 16 07:16:07 UTC 2020 until: Fri Mar 10 07:16:07 UTC 2045
Certificate fingerprints:
        MD5:  ED:AC:F1:7F:88:05:83:2A:83:C0:09:03:D5:00:CA:7B
        SHA1: DC:B5:3F:37:DF:BD:E0:5C:A4:B7:F4:4C:96:12:75:7A:16:C7:61:37
        SHA256:
F2:5B:DE:8A:F2:31:9D:E6:EF:35:F1:30:6F:DA:05:FF:92:B4:15:96:AA:82:67:E3:3C:C1:69:A3
:E5:27:B9:A5
Signature algorithm name: SHA256WITHRSA
Subject Public Key Algorithm: 4096-bit RSA key
Version: 3
```

Examples of successful logging in `/var/log/syslog`:

```
<182>1 2020-03-22T21:54:34.501Z gc3.prome.local NSX 5187 - [nsx@6876
comp="nsx-manager" subcomp="node-mgmt" username="admin" level="INFO"]
Successfully created CA PEM file /config/vmwarensx-node-api/syslog/92a78d8a-
acfd-4515-b05a-2927b70ae920_ca.pem for logging server vrli.prome.local:6514
<182>1 2020-03-22T21:54:36.269Z gc3.prome.local NSX 5187 - [nsx@6876
comp="nsx-manager" subcomp="node-mgmt" username="admin" level="INFO"] Successfully
created client CA PEM file /config/vmwarensx-node-api/syslog/92a78d8a-acfd-4515-
b05a-2927b70ae920_client_ca.pem for logging server vrli.prome.local:6514
<182>1 2020-03-22T21:54:36.495Z gc3.prome.local NSX 5187 - [nsx@6876 comp="nsx-
manager" subcomp="node-mgmt" username="root" level="INFO"] cert issuer = /C=US/
L=California/O=GS/CN=Green IntermediateCertification Authority
<182>1 2020-03-22T21:54:36.514Z gc3.prome.local NSX 5187 - [nsx@6876 comp="nsx-
manager" subcomp="node-mgmt" username="root" level="INFO"] cert subject = /C=US/
ST=California/L=HTG/O=GS/CN=gc.promelocal
<182>1 2020-03-22T21:54:36.539Z gc3.prome.local NSX 5187 - [nsx@6876 comp="nsx-
manager" subcomp="node-mgmt" username="root" level="INFO"] certificate trust check
succeeded. status: 200, result:{'status': 'OK'}
<182>1 2020-03-22T21:54:36.612Z gc3.prome.local NSX 5187 - [nsx@6876 comp="nsx-
manager" subcomp="node-mgmt" username="root" level="INFO"] Certificate already
exists, skip import
<182>1 2020-03-22T21:54:37.322Z gc3.prome.local NSX 5187 - [nsx@6876
comp="nsx-manager" subcomp="node-mgmt" username="admin" level="INFO"] Successfully
created certificate PEM file /config/vmwarensx-node-api/syslog/92a78d8a-acfd-4515-
b05a-2927b70ae920_cert.pem for logging server vrli.prome.local:6514
```

```
<182>1 2020-03-22T21:54:38.020Z gc3.prome.local NSX 5187 - [nsx@6876
comp="nsx-manager" subcomp="node-mgmt" username="admin" level="INFO"] Successfully
created key PEM file /config/vmwarensx-node-api/syslog/92a78d8a-acfd-4515-
b05a-2927b70ae920_key.pem for logging server vrli.prome.local:6514
```

Examples of logging failure in `/var/log/syslog`:

```
<182>1 2020-03-22T21:33:30.424Z gc3.prome.local NSX 5187 - [nsx@6876
comp="nsx-manager" subcomp="node-mgmt" username="admin" level="INFO"]
Successfully created client CA PEM file /config/vmwarensx-node-api/
syslog/76332782-1ec6-483a-95d4-2adeaf2ef112_client_ca.pem for logging server
vrli.prome.local:6514
<182>1 2020-03-22T21:33:30.779Z gc3.prome.local NSX 5187 - [nsx@6876 comp="nsx-
manager" subcomp="node-mgmt" username="root" level="INFO"] cert issuer = /C=US/
L=California/O=GS/CN=Green IntermediateCertification Authority
<182>1 2020-03-22T21:33:30.803Z gc3.prome.local NSX 5187 - [nsx@6876 comp="nsx-
manager" subcomp="node-mgmt" username="root" level="INFO"] cert subject = /C=US/
ST=California/L=HTG/O=GS/CN=gc.promelocal
<179>1 2020-03-22T21:33:30.823Z gc3.prome.local NSX 5187 - [nsx@6876 comp="nsx-
manager" subcomp="node-mgmt" username="root" level="ERROR" errorCode="NODE10"]
Certificate trust check failed. status:200, result: {'error_message': 'Certificate
CN=gc.prome.local,O=GS,L=HTG,ST=California,C=US was not verifiably signed by
CN=gc.prome.local,O=GS,L=HTG,ST=California,C=US: certificate does not verifywith
supplied key', 'status': 'ERROR'}
<179>1 2020-03-22T21:33:30.824Z gc3.prome.local NSX 5187 - [nsx@6876 comp="nsx-
manager" subcomp="node-mgmt" username="admin" level="ERROR" errorCode="NODE10"]
Failed to create certificate PEM file config/vmware/nsx-node-api/
syslog/76332782-1ec6-483a-95d4-2adeaf2ef112_cert.pem for logging server
vrli.prome.local:6514
<182>1 2020-03-22T21:33:31.578Z gc3.prome.local NSX 5187 - [nsx@6876
comp="nsx-manager" subcomp="node-mgmt" username="admin" level="INFO"]
Successfully deleted CA PEM file /config/vmwarensx-node-api/syslog/
76332782-1ec6-483a-95d4-2adeaf2ef112_ca.pem
<182>1 2020-03-22T21:33:32.342Z gc3.prome.local NSX 5187 - [nsx@6876
comp="nsx-manager" subcomp="node-mgmt" username="admin" level="INFO"]
Successfully deleted client CA PEM file /config/vmwarensx-node-api/syslog/
76332782-1ec6-483a-95d4-2adeaf2ef112_ca.pem
<182>1 2020-03-22T21:33:32.346Z gc3.prome.local NSX 16698 - [nsx@6876 comp="nsx-
cli" subcomp="node-mgmt" username="admin" level="INFO" audit="true"] CMD: set
logging-server vrli.prome.local prototls level info serverca Orange-CA.crt.pem
clientca Orange-CA.crt.pem certifi
cate gc-nsxt-mgr.crt.pem key gc-nsxt-mgr.key.pem (duration: 6.365s), Operation
status: CMD_EXECUTED
```

You can check if the certificate and private key match with the following command. If they match, the output will be `writing RSA key`. Any other output means they do not match. For example:

```
root@caserver:~/server-certs# diff  <(openssl x509 -in certs/gc-nsxt-mgr.crt.pem
-pubkey -noout) <(openssl rsa -in private/gc-nsxt-mgr.key.pem -pubout)
writing RSA key
```

Example of a corrupt private key:

```
root@caserver:~/server-certs# diff  <(openssl x509 -in certs/gc-nsxt-mgr.crt.pem
-pubkey -noout) <(openssl rsa -in private/gc-nsxt-mgr-corrupt.key.pem -pubout)
unable to load Private Key
140404188370584:error:0D07209B:asn1 encoding routines:ASN1_get_object:too
long:asn1_lib.c:147:
140404188370584:error:0D068066:asn1 encoding routines:ASN1_CHECK_TLEN:bad object
header:tasn_dec.c:1205:
140404188370584:error:0D07803A:asn1 encoding routines:ASN1_ITEM_EX_D2I:nested asn1
error:tasn_dec.c:386:Type=RSA
140404188370584:error:04093004:rsa routines:OLD_RSA_PRIV_DECODE:RSA
lib:rsa_ameth.c:119:
140404188370584:error:0D07209B:asn1 encoding routines:ASN1_get_object:too
long:asn1_lib.c:147:
140404188370584:error:0D068066:asn1 encoding routines:ASN1_CHECK_TLEN:bad object
header:tasn_dec.c:1205:
140404188370584:error:0D07803A:asn1 encoding routines:ASN1_ITEM_EX_D2I:nested asn1
error:tasn_dec.c:386:Type=PKCS8_PRIV_KEY_INFO
140404188370584:error:0907B00D:PEM routines:PEM_READ_BIO_PRIVATEKEY:ASN1
lib:pem_pkey.c:141:
1,14d0
< -----BEGIN PUBLIC KEY-----
< MIICIjANBgkqhkiG9w0BAQEFAAOCAg8AMIICCgKCAgEAv3yH7pZidfkLrEP3zVa9
< EcOKXlFFjkThZRZMfguenlm8s6QHYVvuUX8IRB48Li3/DUfOj0bzaPWktpv+Q2P0
< N/j4LoX2RzjV/DPxYfLP6GMNMc21L3s9ruBeWUthtUP8khCWd2d2rZ09cUZVl0P9
< kIYBb5RMFC7Z1OUtH3bKdepEf+sXz3DaKZ/WySzYq9x86QDaA3ABO3Q0i7txBscI
< FvXuMDOMQaC3pPp9FWO6IPRAWB57wahLJv6K5qGIfwubSBFg53grT4snf1lDZAhZ
< 9hz5JgGr80GVyWyb7rgigpl9iUWAZx8U9De9XoxmvBN5iEGTIuKGaEgICL176crb
< RMkhjnCqNHI+z6sQvpYJ7U0zZc72eBIWoHUkcWWk3eU6Oy4OiyW6jYuXG7hZYlly
< nSkme3mZUWJKvcoX05+3zeCP623/HzE7X2sNyWFjzeF3XEvauZrIbsJh/xp2ShDa
< uKKEY0gUGhLtCa3TpV9l8d6tFWVy8XjVjdjoVt4s7MfUo/airVmRykfsWrKyNUOQ
< qRZvSbqjt8pm+3bSvKdXX4ul7ptPG2GF20ETWHPwjk2JwQpGhR9zK8fsKzvm6hXi
< kq76zI4FefuVps3e1r39+0F+p6d6i2oUoo24sC1iSePTDhU74efVp6iv8HmnDgYX
< Ylm6Kusr0JT5TJFDfASmrj8CAwEAAQ==
< -----END PUBLIC KEY-----
```

Example of a valid private key and certificate but they are not made for each other:

```
  root@caserver:~/server-certs# diff  <(openssl x509 -in certs/gc-nsxt-mgr.crt.pem
-pubkey -noout) <(openssl rsa -in private/vrli.key.pem -pubout)
  writing RSA key
  2,13c2,13
  < MIICIjANBgkqhkiG9w0BAQEFAAOCAg8AMIICCgKCAgEAv3yH7pZidfkLrEP3zVa9
  < EcOKXlFFjkThZRZMfguenlm8s6QHYVvuUX8IRB48Li3/DUfOj0bzaPWktpv+Q2P0
  < N/j4LoX2RzjV/DPxYfLP6GMNMc21L3s9ruBeWUthtUP8khCWd2d2rZ09cUZVl0P9
  < kIYBb5RMFC7Z1OUtH3bKdepEf+sXz3DaKZ/WySzYq9x86QDaA3ABO3Q0i7txBscI
  < FvXuMDOMQaC3pPp9FWO6IPRAWB57wahLJv6K5qGIfwubSBFg53grT4snf1lDZAhZ
  < 9hz5JgGr80GVyWyb7rgigpl9iUWAZx8U9De9XoxmvBN5iEGTIuKGaEgICL176crb
  < RMkhjnCqNHI+z6sQvpYJ7U0zZc72eBIWoHUkcWWk3eU6Oy4OiyW6jYuXG7hZYlly
  < nSkme3mZUWJKvcoX05+3zeCP623/HzE7X2sNyWFjzeF3XEvauZrIbsJh/xp2ShDa
  < uKKEY0gUGhLtCa3TpV9l8d6tFWVy8XjVjdjoVt4s7MfUo/airVmRykfsWrKyNUOQ
  < qRZvSbqjt8pm+3bSvKdXX4ul7ptPG2GF20ETWHPwjk2JwQpGhR9zK8fsKzvm6hXi
  < kq76zI4FefuVps3e1r39+0F+p6d6i2oUoo24sC1iSePTDhU74efVp6iv8HmnDgYX
  < Ylm6Kusr0JT5TJFDfASmrj8CAwEAAQ==
```

```
---
> MIICIjANBgkqhkiG9w0BAQEFAAOCAg8AMIICCgKCAgEAqvsjay7+o7gCW7szT3ho
> bC34XX2l6u5Jl4/X/pUDI/YHmIf06bsZ1r/l4bTL4Q7BM6+9MI6UYEE7DxUoINGO
> o4FEEQE32KWVFe3gw3homHU39q4pQjsJsxTcTE3oDMlIY0nWJ0PRUst3DffyUH1L
> W0NUN9YdN+fAl2Uf021iuDqVy9V8AH3ON6fu+QCA8nt71ZkzeTxSA0ldpl2NA17F
> rD8rm05wxnV7WtuV7V8PstISiClzhHgZRMl+B0r30OitnyAzEGLaRT3//PKfe0Oe
> HCdxGMlrUtMqxIItJahEsqvMufyqNYecVscyXLHPelizKCsQfy8cO8LnznG8VAdc
> YILSn3uYGZap6aF1SgVxsvZicwvlYnssmgE13Af0nScmfM96k9h5joHVEkWK6O8v
> oT5DGG1kVL2Qly97x0b6EnzUorzivv5zJMKvFcOektR8HdMHQit5uvmMRY3S5zow
> FtvfSDfWxxKyTy6GBrpP+8F+Jq91yGy/qa9lhKBzT2lg+rJp7T8k7/Nm9Tjyx7jL
> EqgEKZEL4chxpo8ucF98hbvXWRuaFHC2iDzGuUmuS1FfjVvHTuIbEMQfjapLZrHx
> 8jHfOP/PL+6kPbvNZZ2rTpczuEoGTQFFW9vX48GzIEyMeR6QWpPR0F7r4xak68P5
> 2PJmMveinDhU35IqWEXHAwcCAwEAAQ==
```

c To view the logging configuration, run the `get logging-server` command. For example,

```
nsx> get logging-servers
192.168.110.60 proto udp level info facility syslog messageid SYSTEM,FABRIC
192.168.110.60 proto udp level info facility auth,user
```

d To clear the remote logging configuration, run the following command:

```
nsx> clear logging-servers
```

2 To configure remote logging on an ESXi host:

a Run the following commands to configure syslog and send a test message:

```
esxcli network firewall ruleset set -r syslog -e true
esxcli system syslog config set --loghost=udp://<log server IP>:<port>
esxcli system syslog reload
esxcli system syslog mark -s "This is a test message"
```

b You can run the following command to display the configuration:

```
esxcli system syslog config get
```

Add Syslog Servers for NSX Nodes

You can use the **Node Profiles** page in NSX Manager to add syslog servers for NSX Manager and NSX Edge nodes.

Procedure

1 From your browser, log in with **admin** privileges to an NSX Manager at `https://`*`nsx-manager-ip-address`*.

2 Navigate to **System > Fabric > Nodes**.

3 Click the **Node Profiles** tab.

4 In the **Syslog Servers** section, click **Add** to add a syslog server.

   a Enter an FQDN or IP address of the syslog server.

   b Specify a port number.

   c Select a protocol.

      The available protocols are **TCP**, **UDP**, and **LI (Aria Operations for Logs)**.

   d Select a log level.

      The available levels are **Emergency**, **Alert**, **Critical**, **Error**, **Warning**, **Notice**, **Information**, and **Debug**.

   e Click **Add**.

5 Repeat step 4 to add more syslog servers, if required.

Log Message IDs

In a log message, the message ID field identifies the type of message. You can use the `messageid` parameter in the `set logging-server` command to filter which log messages are sent to a logging server.

Table 27-7. Log Message IDs

Message ID	Examples
FABRIC	Host node Host preparation Edge node Transport zone Transport node Uplink profiles Cluster profiles Edge cluster
SWITCHING	Logical switch Logical switch ports Switching profiles switch security features
ROUTING	Logical router Logical router ports Static routing Dynamic routing NAT
FIREWALL	Firewall rules Firewall rule sections
FIREWALL-PKTLOG	Firewall connection logs Firewall packet logs

Table 27-7. Log Message IDs (continued)

Message ID	Examples
GROUPING	IP sets Mac sets NSGroups NSServices NSService groups VNI Pool IP Pool
DHCP	DHCP relay
SYSTEM	Appliance management (remote syslog, ntp, etc) Cluster management Trust management Licensing User and roles Task management Install Upgrade (NSX Manager, NSX Edge and host-packages upgrades) Realization Tags
MONITORING	SNMP Port connection Traceflow
-	All other log messages.

Troubleshooting Syslog Issues

If logs are not received by the remote log server, perform the following steps.

- Verify the remote log server's IP address.
- Verify that the `level` parameter is configured correctly.
- Verify that the `facility` parameter is configured correctly.
- If the protocol is TLS, set the protocol to UDP to see if there is a certificate mismatch.
- If the protocol is TLS, verify that port 6514 is open on both ends.
- Remove the message ID filter and see if logs are received by the server.
- Restart the rsyslog service with the command `restart service syslog`.

To learn more about how to configure NSX appliances and hypervisors to send log messages to a remote logging server, see Configure Remote Logging.

Configure Serial Logging on an Appliance VM

You can configure serial logging on an appliance VM to capture log messages when the VM crashes.

Procedure

1. Log in to the VM as `root`.
2. Edit `/etc/default/grub`.
3. Find the parameter `GRUB_CMDLINE_LINUX_DEFAUL` and append `console=ttyS0 console=tty0`.
4. Run the command `update-grub2`.
5. Verify that the `/boot/grub/grub.cfg` file has the change made in step 3.
6. Power off the VM.
7. Edit the VM's configuration (.vmx) file and add the following lines:

```
serial0.present = "TRUE"
serial0.fileType = "file"
serial0.fileName = "serial.out"
serial0.yieldOnMsrRead = "TRUE"
answer.msg.serial.file.open = "Append"
```

8. Power on the VM.

Results

If a kernel panic occurs in the VM, you can find the file `serial.out` containing log messages at the same location as that of the `.vmx` file.

Firewall Audit Log Messages

Firewall configuration changes are audited. Below are examples of audit log messages related to these changes.

Distributed firewall changes in Policy mode

Adding a firewall section (`SecurityPolicy-1`) with a rule (`Rule1_1`):

```
<182>1 2020-08-11T21:58:50.319Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-
manager" level="INFO" reqId="2aff6b4f-3d4f-4d62-a639-61291f7e879e" splitId="a5mxlu78"
splitIndex="1 of 2" subcomp="policy" update="true"] UserName="admin",
ModuleName="DfwSecurityPolicy", Operation="PatchSecurityPolicyForDomain", Operation
status="success", Old
value=[{"precedence":10,"category":"Application","resource_type":"CommunicationMap","id":"Secu
rityPolicy-1","display_name":"SecurityPolicy-1","path":"/infra/domains/default/security-
policies/SecurityPolicy-1","relative_path":"SecurityPolicy-1","parent_path":"/infra/domains/
default","unique_id":"895eeac5-641b-4306-be7f-
a43fdd969ee5","marked_for_delete":false,"overridden":false,"_create_user":"admin","_create_tim
e":1597183130247,"_last_modified_user":"admin","_last_modified_time":1597183130247,"_system_ow
```

ned":false,"_protection":"NOT_PROTECTED","_revision":0}], New value=["default"
"SecurityPolicy-1"
{"resource_type":"SecurityPolicy","id":"SecurityPolicy-1","display_name":"SecurityPolicy-1","p
ath":"/infra/domains/default/security-policies/SecurityPolicy-1","children":[{"Rule":
{"action":"ALLOW","resource_type":"Rule","id":"Rule1_1","display_name":"Rule1_1","path":

<182>1 2020-08-11T21:58:50.320Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-
manager" level="INFO" reqId="2aff6b4f-3d4f-4d62-a639-61291f7e879e" splitId="a5mxlu78"
splitIndex="2 of 2" subcomp="policy" update="true"] "/infra/domains/default/security-policies/
SecurityPolicy-1/rules/
Rule1_1","marked_for_delete":false,"overridden":false,"sequence_number":10,"sources_excluded":
false,"destinations_excluded":false,"source_groups":["ANY"],"destination_groups":
["ANY"],"services":["ANY"],"profiles":["ANY"],"logged":false,"scope":
["ANY"],"disabled":false,"direction":"IN_OUT","ip_protocol":"IPV4_IPV6","_protection":"UNKNOWN
"},"resource_type":"ChildRule","marked_for_delete":false,"mark_for_override":false,"_protectio
n":"UNKNOWN"}],"marked_for_delete":false,"overridden":false,"sequence_number":10,"category":"A
pplication","stateful":true,"locked":false,"scope":["ANY"],"_protection":"UNKNOWN"}]

<182>1 2020-08-11T21:58:50.404Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-
manager" entId="Rule1_1" level="INFO" reqId="2aff6b4f-3d4f-4d62-a639-61291f7e879e"
splitId="E993J2LF" splitIndex="1 of 2" subcomp="policy" update="true"] UserName="admin",
ModuleName="DfwSecurityPolicy", Operation="UpdateSecurityRule", Operation status="success",
Old value=[{"sequence_number":10,"source_groups":["ANY"],"destination_groups":
["ANY"],"services":["ANY"],"action":"ALLOW","logged":false,"scope":
["ANY"],"disabled":false,"direction":"IN_OUT","resource_type":"CommunicationEntry","id":"Rule1
_1","display_name":"Rule1_1","path":"/infra/domains/default/security-policies/
SecurityPolicy-1/rules/Rule1_1","relative_path":"Rule1_1","parent_path":"/infra/domains/
default/security-policies/
SecurityPolicy-1","unique_id":"2024","marked_for_delete":false,"overridden":false,"_create_use
r":"admin","_create_time":1597183130364,"_last_modified_user":"admin","_last_modified_time":15
97183130364,"_system_owned":false,"_protection":"NOT_PROTECTED","_revision":0}], New
value=["default" "SecurityPolicy-1" "Rule1_1"
{"action":"ALLOW","resource_type":"Rule","id":"Rule1_1","display_name":"Rule1_1","path":

<182>1 2020-08-11T21:58:50.404Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-
manager" entId="Rule1_1" level="INFO" reqId="2aff6b4f-3d4f-4d62-a639-61291f7e879e"
splitId="E993J2LF" splitIndex="2 of 2" subcomp="policy" update="true"] "/infra/domains/
default/security-policies/SecurityPolicy-1/rules/
Rule1_1","marked_for_delete":false,"overridden":false,"sequence_number":10,"sources_excluded":
false,"destinations_excluded":false,"source_groups":["ANY"],"destination_groups":
["ANY"],"services":["ANY"],"profiles":["ANY"],"logged":false,"scope":
["ANY"],"disabled":false,"direction":"IN_OUT","ip_protocol":"IPV4_IPV6","_protection":"UNKNOWN
"}]

<182>1 2020-08-11T21:58:50.466Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-
manager" level="INFO" reqId="2aff6b4f-3d4f-4d62-a639-61291f7e879e" splitId="iMHW1shi"
splitIndex="1 of 2" subcomp="policy" update="true"] UserName="admin", ModuleName="Policy",
Operation="PatchInfra", Operation status="success", New
value=[{"enforce_revision_check":true} {"resource_type":"Infra","children":[{"children":
[{"SecurityPolicy":
{"resource_type":"SecurityPolicy","id":"SecurityPolicy-1","display_name":"SecurityPolicy-1","p
ath":"/infra/domains/default/security-policies/SecurityPolicy-1","children":[{"Rule":

```
{"action":"ALLOW","resource_type":"Rule","id":"Rule1_1","display_name":"Rule1_1","path":"/
infra/domains/default/security-policies/SecurityPolicy-1/rules/
Rule1_1","marked_for_delete":false,"overridden":false,"sequence_number":10,"sources_excluded":
false,"destinations_excluded":false,"source_groups":["ANY"],"destination_groups":
["ANY"],"services":["ANY"],"profiles":["ANY"],"logged":false,"scope":
["ANY"],"disabled":false,"direction":"IN_OUT","ip_protocol":"IPV4_IPV6","_protection":"UNKNOWN
"},"resource_type":"ChildRule","marked_for_delete":false,"mark_for_override":false,"_protectio
n":"UNKNOWN"}],"marked_for_delete"

<182>1 2020-08-11T21:58:50.466Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-
manager" level="INFO" reqId="2aff6b4f-3d4f-4d62-a639-61291f7e879e" splitId="iMHW1shi"
splitIndex="2 of 2" subcomp="policy"
update="true"] :false,"overridden":false,"sequence_number":10,"category":"Application","statef
ul":true,"locked":false,"scope":
["ANY"],"_protection":"UNKNOWN"},"resource_type":"ChildSecurityPolicy","marked_for_delete":fal
se,"mark_for_override":false,"_protection":"UNKNOWN"}],"target_type":"Domain","resource_type":
"ChildResourceReference","id":"default","marked_for_delete":false,"mark_for_override":false,"_
protection":"UNKNOWN"}],"marked_for_delete":false,"overridden":false,"_protection":"UNKNOWN","
_revision":-1}]
```


Updating a rule (from `Rule1_1` to `Rule1_1_updated`) in a section (`SecurityPolicy-1`):


```
<182>1 2020-08-11T22:22:06.303Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-
manager" level="INFO" reqId="6aadd8de-d157-4479-b84c-8410dd48c2aa" splitId="mJ7hQGhg"
splitIndex="1 of 2" subcomp="policy" update="true"] UserName="admin",
ModuleName="DfwSecurityPolicy", Operation="PatchSecurityPolicyForDomain", Operation
status="success", New value=["default" "SecurityPolicy-1"
{"resource_type":"SecurityPolicy","id":"SecurityPolicy-1","display_name":"SecurityPolicy-1","p
ath":"/infra/domains/default/security-policies/
SecurityPolicy-1","unique_id":"895eeac5-641b-4306-be7f-a43fdd969ee5","children":[{"Rule":
{"action":"ALLOW","resource_type":"Rule","id":"Rule1_1","display_name":"Rule1_1_updated","path
":"/infra/domains/default/security-policies/SecurityPolicy-1/rules/
Rule1_1","unique_id":"2024","marked_for_delete":false,"overridden":false,"rule_id":2024,"seque
nce_number":10,"sources_excluded":false,"destinations_excluded":false,"source_groups":
["ANY"],"destination_groups":["ANY"],"services":["ANY"],"profiles":
["ANY"],"logged":false,"scope":
["ANY"],"disabled":false,"direction":"IN_OUT","ip_protocol":"IPV4_IPV6","is_default":false,"_p
rotection":"UNKNOWN","_revision":0},"resource_type":
<182>1 2020-08-11T22:22:06.303Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-
manager" level="INFO" reqId="6aadd8de-d157-4479-b84c-8410dd48c2aa" splitId="mJ7hQGhg"
splitIndex="2 of 2" subcomp="policy" update="true"]
"ChildRule","marked_for_delete":false,"mark_for_override":false,"_protection":"UNKNOWN"}],"mar
ked_for_delete":false,"overridden":false,"sequence_number":10,"internal_sequence_number":13000
010,"category":"Application","stateful":true,"tcp_strict":true,"locked":false,"lock_modified_t
ime":0,"scope":["ANY"],"is_default":false,"_protection":"UNKNOWN","_revision":0}]

<182>1 2020-08-11T22:22:06.324Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-
manager" entId="Rule1_1" level="INFO" reqId="6aadd8de-d157-4479-b84c-8410dd48c2aa"
splitId="JKVilI6n" splitIndex="1 of 2" subcomp="policy" update="true"] UserName="admin",
ModuleName="DfwSecurityPolicy", Operation="UpdateSecurityRule", Operation status="success",
Old value=[{"sequence_number":10,"source_groups":["ANY"],"destination_groups":
["ANY"],"services":["ANY"],"action":"ALLOW","logged":false,"scope":
["ANY"],"disabled":false,"direction":"IN_OUT","resource_type":"CommunicationEntry","id":"Rule1
_1","display_name":"Rule1_1","path":"/infra/domains/default/security-policies/
SecurityPolicy-1/rules/Rule1_1","relative_path":"Rule1_1","parent_path":"/infra/domains/
```

default/security-policies/
SecurityPolicy-1","unique_id":"2024","marked_for_delete":false,"overridden":false,"_create_use
r":"admin","_create_time":1597183130364,"_last_modified_user":"admin","_last_modified_time":15
97183130369,"_system_owned":false,"_protection":"NOT_PROTECTED","_revision":0}], New
value=["default" "SecurityPolicy-1" "Rule1_1"
{"action":"ALLOW","resource_type":"Rule","id":"Rule1_1","display_name":"Rule1_1_updated","path
":

<182>1 2020-08-11T22:22:06.324Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-
manager" entId="Rule1_1" level="INFO" reqId="6aadd8de-d157-4479-b84c-8410dd48c2aa"
splitId="JKVilI6n" splitIndex="2 of 2" subcomp="policy" update="true"] "/infra/domains/
default/security-policies/SecurityPolicy-1/rules/
Rule1_1","unique_id":"2024","marked_for_delete":false,"overridden":false,"rule_id":2024,"seque
nce_number":10,"sources_excluded":false,"destinations_excluded":false,"source_groups":
["ANY"],"destination_groups":["ANY"],"services":["ANY"],"profiles":
["ANY"],"logged":false,"scope":
["ANY"],"disabled":false,"direction":"IN_OUT","ip_protocol":"IPV4_IPV6","is_default":false,"_p
rotection":"UNKNOWN","_revision":0}]

<182>1 2020-08-11T22:22:06.363Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-
manager" level="INFO" reqId="6aadd8de-d157-4479-b84c-8410dd48c2aa" splitId="9MtbEpd8"
splitIndex="1 of 2" subcomp="policy" update="true"] UserName="admin", ModuleName="Policy",
Operation="PatchInfra", Operation status="success", New
value=[{"enforce_revision_check":true} {"resource_type":"Infra","children":[{"children":
[{"SecurityPolicy":
{"resource_type":"SecurityPolicy","id":"SecurityPolicy-1","display_name":"SecurityPolicy-1","p
ath":"/infra/domains/default/security-policies/
SecurityPolicy-1","unique_id":"895eeac5-641b-4306-be7f-a43fdd969ee5","children":[{"Rule":
{"action":"ALLOW","resource_type":"Rule","id":"Rule1_1","display_name":"Rule1_1_updated","path
":"/infra/domains/default/security-policies/SecurityPolicy-1/rules/
Rule1_1","unique_id":"2024","marked_for_delete":false,"overridden":false,"rule_id":2024,"seque
nce_number":10,"sources_excluded":false,"destinations_excluded":false,"source_groups":
["ANY"],"destination_groups":["ANY"],"services":["ANY"],"profiles":
["ANY"],"logged":false,"scope":
["ANY"],"disabled":false,"direction":"IN_OUT","ip_protocol":"IPV4_IPV6","is_default":false,"_p
rotection":"UNKNOWN","_revision":

<182>1 2020-08-11T22:22:06.363Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-
manager" level="INFO" reqId="6aadd8de-d157-4479-b84c-8410dd48c2aa" splitId="9MtbEpd8"
splitIndex="2 of 2" subcomp="policy" update="true"]
0},"resource_type":"ChildRule","marked_for_delete":false,"mark_for_override":false,"_protectio
n":"UNKNOWN"}],"marked_for_delete":false,"overridden":false,"sequence_number":10,"internal_seq
uence_number":13000010,"category":"Application","stateful":true,"tcp_strict":true,"locked":fal
se,"lock_modified_time":0,"scope":
["ANY"],"is_default":false,"_protection":"UNKNOWN","_revision":0},"resource_type":"ChildSecuri
tyPolicy","marked_for_delete":false,"mark_for_override":false,"_protection":"UNKNOWN"}],"targe
t_type":"Domain","resource_type":"ChildResourceReference","id":"default","marked_for_delete":f
alse,"mark_for_override":false,"_protection":"UNKNOWN"}],"marked_for_delete":false,"overridden
":false,"_protection":"UNKNOWN","_revision":-1}]

Deleting a rule (`Rule1_2`) from a section (`SecurityPolicy-1`):

```
<182>1 2020-08-11T22:12:24.444Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-
manager" entId="Rule1_2" level="INFO" reqId="1a58e753-460c-443f-8a28-0d40d8af9b76"
subcomp="policy" update="true" username="admin"] UserName="admin",
ModuleName="DfwSecurityPolicy", Operation="DeleteSecurityRule", Operation status="success",
Old value=[{"sequence_number":20,"source_groups":["ANY"],"destination_groups":
["ANY"],"services":["ANY"],"action":"ALLOW","logged":false,"scope":
["ANY"],"disabled":false,"direction":"IN_OUT","resource_type":"CommunicationEntry","id":"Rule1
_2","display_name":"Rule1_2","path":"/infra/domains/default/security-policies/
SecurityPolicy-1/rules/Rule1_2","relative_path":"Rule1_2","parent_path":"/infra/domains/
default/security-policies/
SecurityPolicy-1","unique_id":"2026","marked_for_delete":false,"overridden":false,"_create_use
r":"admin","_create_time":1597183904580,"_last_modified_user":"admin","_last_modified_time":15
97183904582,"_system_owned":false,"_protection":"NOT_PROTECTED","_revision":0}], New
value=["default" "SecurityPolicy-1" "Rule1_2"]

<182>1 2020-08-11T22:12:24.463Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-
manager" level="INFO" reqId="1a58e753-460c-443f-8a28-0d40d8af9b76" splitId="hoDI5YJQ"
splitIndex="1 of 2" subcomp="policy" update="true"] UserName="admin",
ModuleName="DfwSecurityPolicy", Operation="PatchSecurityPolicyForDomain", Operation
status="success", New value=["default" "SecurityPolicy-1"
{"resource_type":"SecurityPolicy","id":"SecurityPolicy-1","display_name":"SecurityPolicy-1","p
ath":"/infra/domains/default/security-policies/
SecurityPolicy-1","unique_id":"895eeac5-641b-4306-be7f-a43fdd969ee5","children":[{"Rule":
{"resource_type":"Rule","id":"Rule1_2","path":"/infra/domains/default/security-policies/
SecurityPolicy-1/rules/
Rule1_2","marked_for_delete":true,"overridden":false,"sources_excluded":false,"destinations_ex
cluded":false,"logged":false,"disabled":false,"direction":"IN_OUT","_protection":"UNKNOWN"},"r
esource_type":"ChildRule","marked_for_delete":true,"mark_for_override":false,"_protection":"UN
KNOWN"}],"marked_for_delete":false,"overridden":false,"sequence_number":10,"internal_sequence_
number":13000010,"category":"Application","stateful":true,"tcp_strict":true,"locked":false,"lo
ck_modified_time":0,"scope":

<182>1 2020-08-11T22:12:24.463Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-
manager" level="INFO" reqId="1a58e753-460c-443f-8a28-0d40d8af9b76" splitId="hoDI5YJQ"
splitIndex="2 of 2" subcomp="policy" update="true"]
["ANY"],"is_default":false,"_protection":"UNKNOWN","_revision":0}]


<182>1 2020-08-11T22:12:24.497Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-
manager" level="INFO" reqId="1a58e753-460c-443f-8a28-0d40d8af9b76" splitId="mxpzQHfF"
splitIndex="1 of 2" subcomp="policy" update="true"] UserName="admin", ModuleName="Policy",
Operation="PatchInfra", Operation status="success", New
value=[{"enforce_revision_check":true} {"resource_type":"Infra","children":[{"children":
[{"SecurityPolicy":
{"resource_type":"SecurityPolicy","id":"SecurityPolicy-1","display_name":"SecurityPolicy-1","p
ath":"/infra/domains/default/security-policies/
SecurityPolicy-1","unique_id":"895eeac5-641b-4306-be7f-a43fdd969ee5","children":[{"Rule":
{"resource_type":"Rule","id":"Rule1_2","path":"/infra/domains/default/security-policies/
SecurityPolicy-1/rules/
Rule1_2","marked_for_delete":true,"overridden":false,"sources_excluded":false,"destinations_ex
cluded":false,"logged":false,"disabled":false,"direction":"IN_OUT","_protection":"UNKNOWN"},"r
esource_type":"ChildRule","marked_for_delete":true,"mark_for_override":false,"_protection":"UN
KNOWN"}],"marked_for_delete":false,"overridden":false,"sequence_number":10,"internal_sequence_
```

```
number":13000010,"category":"Application","stateful":true,"tcp_strict":
<182>1 2020-08-11T22:12:24.497Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-
manager" level="INFO" reqId="1a58e753-460c-443f-8a28-0d40d8af9b76" splitId="mxpzQHfF"
splitIndex="2 of 2" subcomp="policy" update="true"]
true,"locked":false,"lock_modified_time":0,"scope":
["ANY"],"is_default":false,"_protection":"UNKNOWN","_revision":0},"resource_type":"ChildSecuri
tyPolicy","marked_for_delete":false,"mark_for_override":false,"_protection":"UNKNOWN"}],"targe
t_type":"Domain","resource_type":"ChildResourceReference","id":"default","marked_for_delete":f
alse,"mark_for_override":false,"_protection":"UNKNOWN"}],"marked_for_delete":false,"overridden
":false,"_protection":"UNKNOWN","_revision":-1}]
```

Deleting a section (`SecurityPolicy-1`) that contains a rule (`Rule1_1`):

```
<182>1 2020-08-11T22:24:24.898Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-
manager" entId="SecurityPolicy-1" level="INFO" reqId="724b5494-10cd-4124-a431-56ba7d922bbf"
splitId="4WIXz9qL" splitIndex="1 of 2" subcomp="policy" update="true" username="admin"]
UserName="admin", ModuleName="DfwSecurityPolicy", Operation="DeleteSecurityPolicyForDomain",
Operation status="success", Old
value=[{"precedence":10,"category":"Application","resource_type":"CommunicationMap","id":"Secu
rityPolicy-1","display_name":"SecurityPolicy-1","path":"/infra/domains/default/security-
policies/SecurityPolicy-1","relative_path":"SecurityPolicy-1","parent_path":"/infra/domains/
default","unique_id":"895eeac5-641b-4306-be7f-
a43fdd969ee5","marked_for_delete":false,"overridden":false,"_create_user":"admin","_create_tim
e":1597183130247,"_last_modified_user":"admin","_last_modified_time":1597183130251,"_system_ow
ned":false,"_protection":"NOT_PROTECTED","_revision":0}{"sequence_number":10,"source_groups":
["ANY"],"destination_groups":["ANY"],"services":
["ANY"],"action":"ALLOW","logged":false,"scope":
["ANY"],"disabled":false,"direction":"IN_OUT","resource_type":"CommunicationEntry","id":"Rule1
_1","display_name":"Rule1_1_updated","path":

<182>1 2020-08-11T22:24:24.898Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-
manager" entId="SecurityPolicy-1" level="INFO" reqId="724b5494-10cd-4124-a431-56ba7d922bbf"
splitId="4WIXz9qL" splitIndex="2 of 2" subcomp="policy" update="true" username="admin"] "/
infra/domains/default/security-policies/SecurityPolicy-1/rules/
Rule1_1","relative_path":"Rule1_1","parent_path":"/infra/domains/default/security-policies/
SecurityPolicy-1","unique_id":"2024","marked_for_delete":false,"overridden":false,"_create_use
r":"admin","_create_time":1597183130364,"_last_modified_user":"admin","_last_modified_time":15
97184526313,"_system_owned":false,"_protection":"NOT_PROTECTED","_revision":1}], New
value=["default" "SecurityPolicy-1"]

<182>1 2020-08-11T22:24:24.938Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-
manager" level="INFO" reqId="724b5494-10cd-4124-a431-56ba7d922bbf" subcomp="policy"
update="true"] UserName="admin", ModuleName="Policy", Operation="PatchInfra", Operation
status="success", New value=[{"enforce_revision_check":true}
{"resource_type":"Infra","children":[{"children":[{"SecurityPolicy":
{"resource_type":"SecurityPolicy","id":"SecurityPolicy-1","path":"/infra/domains/default/
security-policies/
SecurityPolicy-1","marked_for_delete":true,"overridden":false,"locked":false,"_protection":"UN
KNOWN"},"resource_type":"ChildSecurityPolicy","marked_for_delete":true,"mark_for_override":fal
se,"_protection":"UNKNOWN"}],"target_type":"Domain","resource_type":"ChildResourceReference","
id":"default","marked_for_delete":false,"mark_for_override":false,"_protection":"UNKNOWN"}],"m
arked_for_delete":false,"overridden":false,"_protection":"UNKNOWN","_revision":-1}]
```

Gateway firewall changes in Policy mode

Note that log messages for a tier-0 gateway and a tier-1 gateway are similar.

Adding a section (`T1-Policies`) with a rule (`myT1_Rule1`) for a tier-1 gateway (`myT1`):

```
<182>1 2020-08-11T22:31:26.800Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-
manager" level="INFO" reqId="c2790fbd-db29-46d3-9a0e-1003455ee9ea" splitId="Ta8faYzQ"
splitIndex="1 of 2" subcomp="policy" update="true"] UserName="admin",
ModuleName="PolicyEdgeFirewall", Operation="PatchGatewayPolicyForDomain", Operation
status="success", Old
value=[{"precedence":10,"category":"LocalGatewayRules","resource_type":"CommunicationMap","id"
:"T1-Policies","display_name":"T1-Policies","path":"/infra/domains/default/gateway-
policies/T1-Policies","relative_path":"T1-Policies","parent_path":"/infra/domains/
default","unique_id":"a73c1345-6b4e-43e0-
b4ee-9a91c7ba9df6","marked_for_delete":false,"overridden":false,"_create_user":"admin","_creat
e_time":1597185086789,"_last_modified_user":"admin","_last_modified_time":1597185086789,"_syst
em_owned":false,"_protection":"NOT_PROTECTED","_revision":0}], New value=["default" "T1-
Policies" {"resource_type":"GatewayPolicy","id":"T1-Policies","display_name":"T1-
Policies","path":"/infra/domains/default/gateway-policies/T1-Policies","children":[{"Rule":
{"action":"ALLOW","resource_type":"Rule","id":"myT1_Rule1","display_name":"myT1_Rule1","path":

<182>1 2020-08-11T22:31:26.801Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-
manager" level="INFO" reqId="c2790fbd-db29-46d3-9a0e-1003455ee9ea" splitId="Ta8faYzQ"
splitIndex="2 of 2" subcomp="policy" update="true"] "/infra/domains/default/gateway-
policies/T1-Policies/rules/
myT1_Rule1","marked_for_delete":false,"overridden":false,"sequence_number":10,"sources_exclude
d":false,"destinations_excluded":false,"source_groups":["ANY"],"destination_groups":
["ANY"],"services":["ANY"],"profiles":["ANY"],"logged":false,"scope":["/infra/tier-1s/
myT1"],"disabled":false,"direction":"IN_OUT","ip_protocol":"IPV4_IPV6","_protection":"UNKNOWN"
},"resource_type":"ChildRule","marked_for_delete":false,"mark_for_override":false,"_protection
":"UNKNOWN"}],"marked_for_delete":false,"overridden":false,"sequence_number":10,"category":"Lo
calGatewayRules","stateful":true,"locked":false,"_protection":"UNKNOWN"}]

<182>1 2020-08-11T22:31:26.878Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-
manager" level="INFO" reqId="c2790fbd-db29-46d3-9a0e-1003455ee9ea" splitId="aZfgiFKt"
splitIndex="1 of 2" subcomp="policy" update="true" username="admin"] UserName="admin",
ModuleName="PolicyEdgeFirewall", Operation="PatchGatewayRule", Operation status="success",
Old value=[{"sequence_number":10,"source_groups":["ANY"],"destination_groups":
["ANY"],"services":["ANY"],"action":"ALLOW","logged":false,"scope":["/infra/tier-1s/
myT1"],"disabled":false,"direction":"IN_OUT","resource_type":"CommunicationEntry","id":"myT1_R
ule1","display_name":"myT1_Rule1","path":"/infra/domains/default/gateway-policies/T1-Policies/
rules/myT1_Rule1","relative_path":"myT1_Rule1","parent_path":"/infra/domains/default/gateway-
policies/T1-
Policies","unique_id":"2028","marked_for_delete":false,"overridden":false,"_create_user":"admi
n","_create_time":1597185086809,"_last_modified_user":"admin","_last_modified_time":1597185086
809,"_system_owned":false,"_protection":"NOT_PROTECTED","_revision":0}], New value=["default"
"T1-Policies" "myT1_Rule1"
{"action":"ALLOW","resource_type":"Rule","id":"myT1_Rule1","display_name":"myT1_Rule1","path":

<182>1 2020-08-11T22:31:26.878Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-
manager" level="INFO" reqId="c2790fbd-db29-46d3-9a0e-1003455ee9ea" splitId="aZfgiFKt"
splitIndex="2 of 2" subcomp="policy" update="true" username="admin"] "/infra/domains/default/
gateway-policies/T1-Policies/rules/
myT1_Rule1","marked_for_delete":false,"overridden":false,"sequence_number":10,"sources_exclude
```

```
d":false,"destinations_excluded":false,"source_groups":["ANY"],"destination_groups":
["ANY"],"services":["ANY"],"profiles":["ANY"],"logged":false,"scope":["/infra/tier-1s/
myT1"],"disabled":false,"direction":"IN_OUT","ip_protocol":"IPV4_IPV6","_protection":"UNKNOWN"
}]

<182>1 2020-08-11T22:31:26.890Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-
manager" level="INFO" reqId="c2790fbd-db29-46d3-9a0e-1003455ee9ea" splitId="0s7tdCjN"
splitIndex="1 of 2" subcomp="policy" update="true"] UserName="admin", ModuleName="Policy",
Operation="PatchInfra", Operation status="success", New
value=[{"enforce_revision_check":true} {"resource_type":"Infra","children":[{"children":
[{"GatewayPolicy":{"resource_type":"GatewayPolicy","id":"T1-Policies","display_name":"T1-
Policies","path":"/infra/domains/default/gateway-policies/T1-Policies","children":[{"Rule":
{"action":"ALLOW","resource_type":"Rule","id":"myT1_Rule1","display_name":"myT1_Rule1","path":
"/infra/domains/default/gateway-policies/T1-Policies/rules/
myT1_Rule1","marked_for_delete":false,"overridden":false,"sequence_number":10,"sources_exclude
d":false,"destinations_excluded":false,"source_groups":["ANY"],"destination_groups":
["ANY"],"services":["ANY"],"profiles":["ANY"],"logged":false,"scope":["/infra/tier-1s/
myT1"],"disabled":false,"direction":"IN_OUT","ip_protocol":"IPV4_IPV6","_protection":"UNKNOWN"
},"resource_type":"ChildRule","marked_for_delete":false,"mark_for_override":false,"_protection
":

<182>1 2020-08-11T22:31:26.890Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-
manager" level="INFO" reqId="c2790fbd-db29-46d3-9a0e-1003455ee9ea" splitId="0s7tdCjN"
splitIndex="2 of 2" subcomp="policy" update="true"]
"UNKNOWN"}],"marked_for_delete":false,"overridden":false,"sequence_number":10,"category":"Loca
lGatewayRules","stateful":true,"locked":false,"_protection":"UNKNOWN"},"resource_type":"ChildG
atewayPolicy","marked_for_delete":false,"mark_for_override":false,"_protection":"UNKNOWN"}],"t
arget_type":"Domain","resource_type":"ChildResourceReference","id":"default","marked_for_delet
e":false,"mark_for_override":false,"_protection":"UNKNOWN"}],"marked_for_delete":false,"overri
dden":false,"_protection":"UNKNOWN","_revision":-1}]
```


Updating a rule (from `myT1_Rule1` to `myT1_Rule1_Updated`) in a section (`T1-Policies`):


```
<182>1 2020-08-11T22:36:19.410Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-
manager" level="INFO" reqId="a17fcbdc-1aed-4526-93e9-40a3730eeb7f" splitId="BiHDjsY8"
splitIndex="1 of 2" subcomp="policy" update="true"] UserName="admin",
ModuleName="PolicyEdgeFirewall", Operation="PatchGatewayPolicyForDomain", Operation
status="success", New value=["default" "T1-Policies"
{"resource_type":"GatewayPolicy","id":"T1-Policies","display_name":"T1-Policies","path":"/
infra/domains/default/gateway-policies/T1-Policies","unique_id":"a73c1345-6b4e-43e0-
b4ee-9a91c7ba9df6","children":[{"Rule":
{"action":"ALLOW","resource_type":"Rule","id":"myT1_Rule1","display_name":"myT1_Rule1_Updated"
,"path":"/infra/domains/default/gateway-policies/T1-Policies/rules/
myT1_Rule1","unique_id":"2028","marked_for_delete":false,"overridden":false,"rule_id":2028,"se
quence_number":10,"sources_excluded":false,"destinations_excluded":false,"source_groups":
["ANY"],"destination_groups":["ANY"],"services":["ANY"],"profiles":
["ANY"],"logged":false,"scope":["/infra/tier-1s/
myT1"],"disabled":false,"direction":"IN_OUT","ip_protocol":"IPV4_IPV6","is_default":false,"_pr
otection":"UNKNOWN","_revision":0},"resource_type":

<182>1 2020-08-11T22:36:19.410Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-
manager" level="INFO" reqId="a17fcbdc-1aed-4526-93e9-40a3730eeb7f" splitId="BiHDjsY8"
splitIndex="2 of 2" subcomp="policy" update="true"]
"ChildRule","marked_for_delete":false,"mark_for_override":false,"_protection":"UNKNOWN"}],"mar
ked_for_delete":false,"overridden":false,"sequence_number":10,"internal_sequence_number":13000
```

010,"category":"LocalGatewayRules","stateful":true,"tcp_strict":true,"locked":false,"lock_modified_time":0,"is_default":false,"_protection":"UNKNOWN","_revision":0}]

<182>1 2020-08-11T22:36:19.430Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-manager" level="INFO" reqId="a17fcbdc-1aed-4526-93e9-40a3730eeb7f" splitId="HqttDMqz" splitIndex="1 of 2" subcomp="policy" update="true" username="admin"] UserName="admin", ModuleName="PolicyEdgeFirewall", Operation="PatchGatewayRule", Operation status="success", Old value=[{"sequence_number":10,"source_groups":["ANY"],"destination_groups":["ANY"],"services":["ANY"],"action":"ALLOW","logged":false,"scope":["/infra/tier-1s/myT1"],"disabled":false,"direction":"IN_OUT","resource_type":"CommunicationEntry","id":"myT1_Rule1","display_name":"myT1_Rule1","path":"/infra/domains/default/gateway-policies/T1-Policies/rules/myT1_Rule1","relative_path":"myT1_Rule1","parent_path":"/infra/domains/default/gateway-policies/T1-Policies","unique_id":"2028","marked_for_delete":false,"overridden":false,"_create_user":"admin","_create_time":1597185086809,"_last_modified_user":"admin","_last_modified_time":1597185086841,"_system_owned":false,"_protection":"NOT_PROTECTED","_revision":0}], New value=["default" "T1-Policies" "myT1_Rule1"
{"action":"ALLOW","resource_type":"Rule","id":"myT1_Rule1","display_name":"myT1_Rule1_Updated","path":

<182>1 2020-08-11T22:36:19.430Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-manager" level="INFO" reqId="a17fcbdc-1aed-4526-93e9-40a3730eeb7f" splitId="HqttDMqz" splitIndex="2 of 2" subcomp="policy" update="true" username="admin"] "/infra/domains/default/gateway-policies/T1-Policies/rules/myT1_Rule1","unique_id":"2028","marked_for_delete":false,"overridden":false,"rule_id":2028,"sequence_number":10,"sources_excluded":false,"destinations_excluded":false,"source_groups":["ANY"],"destination_groups":["ANY"],"services":["ANY"],"profiles":["ANY"],"logged":false,"scope":["/infra/tier-1s/myT1"],"disabled":false,"direction":"IN_OUT","ip_protocol":"IPV4_IPV6","is_default":false,"_protection":"UNKNOWN","_revision":0}]

<182>1 2020-08-11T22:36:19.443Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-manager" level="INFO" reqId="a17fcbdc-1aed-4526-93e9-40a3730eeb7f" splitId="fMYsYjV5" splitIndex="1 of 2" subcomp="policy" update="true"] UserName="admin", ModuleName="Policy", Operation="PatchInfra", Operation status="success", New value=[{"enforce_revision_check":true} {"resource_type":"Infra","children":[{"children":[{"GatewayPolicy":{"resource_type":"GatewayPolicy","id":"T1-Policies","display_name":"T1-Policies","path":"/infra/domains/default/gateway-policies/T1-Policies","unique_id":"a73c1345-6b4e-43e0-b4ee-9a91c7ba9df6","children":[{"Rule":{"action":"ALLOW","resource_type":"Rule","id":"myT1_Rule1","display_name":"myT1_Rule1_Updated","path":"/infra/domains/default/gateway-policies/T1-Policies/rules/myT1_Rule1","unique_id":"2028","marked_for_delete":false,"overridden":false,"rule_id":2028,"sequence_number":10,"sources_excluded":false,"destinations_excluded":false,"source_groups":["ANY"],"destination_groups":["ANY"],"services":["ANY"],"profiles":["ANY"],"logged":false,"scope":["/infra/tier-1s/myT1"],"disabled":false,"direction":"IN_OUT","ip_protocol":"IPV4_IPV6","is_default":false,"_protection":"UNKNOWN","_revision":

<182>1 2020-08-11T22:36:19.443Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-manager" level="INFO" reqId="a17fcbdc-1aed-4526-93e9-40a3730eeb7f" splitId="fMYsYjV5" splitIndex="2 of 2" subcomp="policy" update="true"]
0},"resource_type":"ChildRule","marked_for_delete":false,"mark_for_override":false,"_protection":"UNKNOWN"}],"marked_for_delete":false,"overridden":false,"sequence_number":10,"internal_sequence_number":13000010,"category":"LocalGatewayRules","stateful":true,"tcp_strict":true,"locked":false,"lock_modified_time":0,"is_default":false,"_protection":"UNKNOWN","_revision":0},"res

ource_type":"ChildGatewayPolicy","marked_for_delete":false,"mark_for_override":false,"_protect
ion":"UNKNOWN"}],"target_type":"Domain","resource_type":"ChildResourceReference","id":"default
","marked_for_delete":false,"mark_for_override":false,"_protection":"UNKNOWN"}],"marked_for_de
lete":false,"overridden":false,"_protection":"UNKNOWN","_revision":-1}]

Delete a rule (`myT1_Rule2`) from a section (`T1-Policies`):

<182>1 2020-08-11T22:38:03.262Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-
manager" entId="myT1_Rule2" level="INFO" reqId="ccb8d0bb-0fe2-415a-9979-a1a3a80a7038"
subcomp="policy" update="true" username="admin"] UserName="admin",
ModuleName="PolicyEdgeFirewall", Operation="DeleteGatewayRule", Operation status="success",
Old value=[{"sequence_number":20,"source_groups":["ANY"],"destination_groups":
["ANY"],"services":["ANY"],"action":"ALLOW","logged":false,"scope":["/infra/tier-1s/
myT1"],"disabled":false,"direction":"IN_OUT","resource_type":"CommunicationEntry","id":"myT1_R
ule2","display_name":"myT1_Rule2","path":"/infra/domains/default/gateway-policies/T1-Policies/
rules/myT1_Rule2","relative_path":"myT1_Rule2","parent_path":"/infra/domains/default/gateway-
policies/T1-
Policies","unique_id":"2029","marked_for_delete":false,"overridden":false,"_create_user":"admi
n","_create_time":1597185467310,"_last_modified_user":"admin","_last_modified_time":1597185467
314,"_system_owned":false,"_protection":"NOT_PROTECTED","_revision":0}], New value=["default"
"T1-Policies" "myT1_Rule2"]

<182>1 2020-08-11T22:38:03.280Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-
manager" level="INFO" reqId="ccb8d0bb-0fe2-415a-9979-a1a3a80a7038" splitId="GlUhKvqu"
splitIndex="1 of 2" subcomp="policy" update="true"] UserName="admin",
ModuleName="PolicyEdgeFirewall", Operation="PatchGatewayPolicyForDomain", Operation
status="success", New value=["default" "T1-Policies"
{"resource_type":"GatewayPolicy","id":"T1-Policies","display_name":"T1-Policies","path":"/
infra/domains/default/gateway-policies/T1-Policies","unique_id":"a73c1345-6b4e-43e0-
b4ee-9a91c7ba9df6","children":[{"Rule":{"resource_type":"Rule","id":"myT1_Rule2","path":"/
infra/domains/default/gateway-policies/T1-Policies/rules/
myT1_Rule2","marked_for_delete":true,"overridden":false,"sources_excluded":false,"destinations
_excluded":false,"logged":false,"disabled":false,"direction":"IN_OUT","_protection":"UNKNOWN"}
,"resource_type":"ChildRule","marked_for_delete":true,"mark_for_override":false,"_protection":
"UNKNOWN"}],"marked_for_delete":false,"overridden":false,"sequence_number":10,"internal_sequen
ce_number":13000010,"category":"LocalGatewayRules","stateful":true,"tcp_strict":true,"locked":
false,"lock_modified_time":0,"is_default":

<182>1 2020-08-11T22:38:03.280Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-
manager" level="INFO" reqId="ccb8d0bb-0fe2-415a-9979-a1a3a80a7038" splitId="GlUhKvqu"
splitIndex="2 of 2" subcomp="policy" update="true"]
false,"_protection":"UNKNOWN","_revision":0}]

<182>1 2020-08-11T22:38:03.295Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-
manager" level="INFO" reqId="ccb8d0bb-0fe2-415a-9979-a1a3a80a7038" splitId="xnO9T8NE"
splitIndex="1 of 2" subcomp="policy" update="true"] UserName="admin", ModuleName="Policy",
Operation="PatchInfra", Operation status="success", New
value=[{"enforce_revision_check":true} {"resource_type":"Infra","children":[{"children":
[{"GatewayPolicy":{"resource_type":"GatewayPolicy","id":"T1-Policies","display_name":"T1-
Policies","path":"/infra/domains/default/gateway-policies/T1-
Policies","unique_id":"a73c1345-6b4e-43e0-b4ee-9a91c7ba9df6","children":[{"Rule":
{"resource_type":"Rule","id":"myT1_Rule2","path":"/infra/domains/default/gateway-policies/T1-
Policies/rules/
myT1_Rule2","marked_for_delete":true,"overridden":false,"sources_excluded":false,"destinations
_excluded":false,"logged":false,"disabled":false,"direction":"IN_OUT","_protection":"UNKNOWN"}

```
,"resource_type":"ChildRule","marked_for_delete":true,"mark_for_override":false,"_protection":
"UNKNOWN"}],"marked_for_delete":false,"overridden":false,"sequence_number":10,"internal_sequen
ce_number":13000010,"category":"LocalGatewayRules","stateful":true,"tcp_strict":true,"locked":

<182>1 2020-08-11T22:38:03.295Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-
manager" level="INFO" reqId="ccb8d0bb-0fe2-415a-9979-a1a3a80a7038" splitId="xnO9T8NE"
splitIndex="2 of 2" subcomp="policy" update="true"]
false,"lock_modified_time":0,"is_default":false,"_protection":"UNKNOWN","_revision":0},"resour
ce_type":"ChildGatewayPolicy","marked_for_delete":false,"mark_for_override":false,"_protection
":"UNKNOWN"}],"target_type":"Domain","resource_type":"ChildResourceReference","id":"default","
marked_for_delete":false,"mark_for_override":false,"_protection":"UNKNOWN"}],"marked_for_delet
e":false,"overridden":false,"_protection":"UNKNOWN","_revision":-1}]
```

Deleting a section (`T1-Policies`) that contains a rule (`myT1_Rule1_Updated`):

```
<182>1 2020-08-11T22:41:30.726Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-
manager" entId="T1-Policies" level="INFO" reqId="d751343c-32ab-46ee-b176-752b8ae1ec0d"
splitId="WZc3oxDG" splitIndex="1 of 2" subcomp="policy" update="true" username="admin"]
UserName="admin", ModuleName="PolicyEdgeFirewall", Operation="DeleteGatewayPolicy", Operation
status="success", Old
value=[{"precedence":10,"category":"LocalGatewayRules","resource_type":"CommunicationMap","id"
:"T1-Policies","display_name":"T1-Policies","path":"/infra/domains/default/gateway-
policies/T1-Policies","relative_path":"T1-Policies","parent_path":"/infra/domains/
default","unique_id":"a73c1345-6b4e-43e0-
b4ee-9a91c7ba9df6","marked_for_delete":false,"overridden":false,"_create_user":"admin","_creat
e_time":1597185086789,"_last_modified_user":"admin","_last_modified_time":1597185086790,"_syst
em_owned":false,"_protection":"NOT_PROTECTED","_revision":0}
{"sequence_number":10,"source_groups":["ANY"],"destination_groups":["ANY"],"services":
["ANY"],"action":"ALLOW","logged":false,"scope":["/infra/tier-1s/
myT1"],"disabled":false,"direction":"IN_OUT","resource_type":"CommunicationEntry","id":"myT1_R
ule1","display_name":"myT1_Rule1_Updated","path":

<182>1 2020-08-11T22:41:30.726Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-
manager" entId="T1-Policies" level="INFO" reqId="d751343c-32ab-46ee-b176-752b8ae1ec0d"
splitId="WZc3oxDG" splitIndex="2 of 2" subcomp="policy" update="true" username="admin"] "/
infra/domains/default/gateway-policies/T1-Policies/rules/
myT1_Rule1","relative_path":"myT1_Rule1","parent_path":"/infra/domains/default/gateway-
policies/T1-
Policies","unique_id":"2028","marked_for_delete":false,"overridden":false,"_create_user":"admi
n","_create_time":1597185086809,"_last_modified_user":"admin","_last_modified_time":1597185379
419,"_system_owned":false,"_protection":"NOT_PROTECTED","_revision":1}], New value=["default"
"T1-Policies"]

<182>1 2020-08-11T22:41:30.733Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-
manager" level="INFO" reqId="d751343c-32ab-46ee-b176-752b8ae1ec0d" subcomp="policy"
update="true"] UserName="admin", ModuleName="Policy", Operation="PatchInfra", Operation
status="success", New value=[{"enforce_revision_check":true}
{"resource_type":"Infra","children":[{"children":[{"GatewayPolicy":
{"resource_type":"GatewayPolicy","id":"T1-Policies","path":"/infra/domains/default/gateway-
policies/T1-
Policies","marked_for_delete":true,"overridden":false,"locked":false,"_protection":"UNKNOWN"},
```

"resource_type":"ChildGatewayPolicy","marked_for_delete":true,"mark_for_override":false,"_protection":"UNKNOWN"}],"target_type":"Domain","resource_type":"ChildResourceReference","id":"default","marked_for_delete":false,"mark_for_override":false,"_protection":"UNKNOWN"}],"marked_for_delete":false,"overridden":false,"_protection":"UNKNOWN","_revision":-1}]

Distributed firewall changes in Manager mode

Adding a firewall section (`FirewallSection-2`):

```
<182>1 2020-08-12T00:25:53.300Z manager1 NSX 1503 - [nsx@6876 audit="true" comp="nsx-manager"
level="INFO" reqId="244e8a97-93d4-4047-b817-81b59b94ce13" subcomp="manager" username="admin"]
UserName="admin", ModuleName="NSX-Firewall", Operation="CREATE", Operation status="success",
New value=[FirewallSectionLock [Id=0ffb0688-9f4e-4096-a19f-2d98ce8cfbeb,
sectionId=f5226cab-525b-4e33-a26d-e5053fbba0a1, sectionRevision=0, locked=false,
comments=Default section unlock comment, created_by=admin, create_time=1597191953299,
last_modified_by=admin, last_modified_time=1597191953299]]

<182>1 2020-08-12T00:25:53.313Z manager1 NSX 1503 FIREWALL [nsx@6876 audit="true" comp="nsx-
manager" entId="f5226cab-525b-4e33-a26d-e5053fbba0a1" level="INFO" reqId="244e8a97-93d4-4047-
b817-81b59b94ce13" subcomp="manager" update="true" username="admin"] UserName="admin",
ModuleName="Firewall", Operation="AddSection", Operation status="success", New
value=[{"operation":"insert_before","id":"ffffffff-8a04-4924-a5b4-54d30e81befe"}
{"locked":false,"autoplumbed":false,"tcp_strict":false,"display_name":"FirewallSection-2","sec
tion_type":"LAYER3","stateful":true,"_protection":"UNKNOWN"}]
```

Adding a rule (`mp_Rule1`) to a section (`FirewallSection-2`):

```
<182>1 2020-08-12T00:27:21.252Z manager1 NSX 1503 FIREWALL [nsx@6876 audit="true" comp="nsx-
manager" entId="f5226cab-525b-4e33-a26d-e5053fbba0a1" level="INFO" reqId="3562bb3e-bf18-4aa0-
a1dd-abde13e8559c" splitId="ScK9FB8V" splitIndex="1 of 2" subcomp="manager" update="true"
username="admin"] UserName="admin", ModuleName="Firewall",
Operation="UpdateSectionWithRules", Operation status="success", Old
value=[{"locked":false,"comments":"Default section unlock
comment","lock_modified_by":"admin","lock_modified_time":1597191953299,"autoplumbed":false,"en
forced_on":"VIF","tcp_strict":false,"category":"Default","resource_type":"FirewallSection","id
":"f5226cab-525b-4e33-a26d-
e5053fbba0a1","display_name":"FirewallSection-2","section_type":"LAYER3","stateful":true,"rule
_count":0,"is_default":false,"_create_user":"admin","_create_time":1597191953297,"_last_modifi
ed_user":"admin","_last_modified_time":1597191953297,"_system_owned":false,"_protection":"NOT_
PROTECTED","_revision":0}], New value=["f5226cab-525b-4e33-a26d-e5053fbba0a1" {"rules":
[{"display_name":"mp_Rule1","sources_excluded":false,"destinations_excluded":false,"action":"A
LLOW","disabled":false,"logged":false,"direction":"IN_OUT","ip_protocol":"IPV4_IPV6","is_defau
lt":false}],"resource_type":"FirewallSection","id":
<182>1 2020-08-12T00:27:21.252Z manager1 NSX 1503 FIREWALL [nsx@6876 audit="true" comp="nsx-
manager" entId="f5226cab-525b-4e33-a26d-e5053fbba0a1" level="INFO" reqId="3562bb3e-bf18-4aa0-
a1dd-abde13e8559c" splitId="ScK9FB8V" splitIndex="2 of 2" subcomp="manager" update="true"
username="admin"] "f5226cab-525b-4e33-a26d-
e5053fbba0a1","display_name":"FirewallSection-2","section_type":"LAYER3","stateful":true,"rule
_count":0,"is_default":false,"locked":false,"comments":"Default section unlock
comment","lock_modified_by":"admin","lock_modified_time":1597191953299,"autoplumbed":false,"en
forced_on":"VIF","tcp_strict":false,"category":"Default","_protection":"UNKNOWN","_revision":0
}]
```

Updating a rule (from `mp_Rule1` to `mp_Rule1_updated`) in a section (`FirewallSection-2`):

```
<182>1 2020-08-12T00:28:54.226Z manager1 NSX 1503 FIREWALL [nsx@6876 audit="true" comp="nsx-
manager" entId="f5226cab-525b-4e33-a26d-e5053fbba0a1" level="INFO"
reqId="37954994-8d59-448e-923d-940813087640" splitId="KcUAlRYl" splitIndex="1 of 2"
subcomp="manager" update="true" username="admin"] UserName="admin", ModuleName="Firewall",
Operation="UpdateSectionWithRules", Operation status="success", Old
value=[{"section_id":"f5226cab-525b-4e33-a26d-
e5053fbba0a1","resource_type":"FirewallRule","id":"536870917","display_name":"mp_Rule1","sourc
es_excluded":false,"destinations_excluded":false,"action":"ALLOW","disabled":false,"logged":fa
lse,"direction":"IN_OUT","ip_protocol":"IPV4_IPV6","is_default":false}
{"locked":false,"comments":"Default section unlock
comment","lock_modified_by":"admin","lock_modified_time":1597191953299,"autoplumbed":false,"en
forced_on":"VIF","tcp_strict":false,"category":"Default","resource_type":"FirewallSection","id
":"f5226cab-525b-4e33-a26d-
e5053fbba0a1","display_name":"FirewallSection-2","section_type":"LAYER3","stateful":true,"rule
_count":1,"is_default":false,"_create_user":"admin","_create_time":1597191953297,"_last_modifi
ed_user":"admin","_last_modified_time":1597192041235,"_system_owned":false,"_protection":"NOT_
PROTECTED","_revision":
<182>1 2020-08-12T00:28:54.226Z manager1 NSX 1503 FIREWALL [nsx@6876 audit="true" comp="nsx-
manager" entId="f5226cab-525b-4e33-a26d-e5053fbba0a1" level="INFO"
reqId="37954994-8d59-448e-923d-940813087640" splitId="KcUAlRYl" splitIndex="2 of 2"
subcomp="manager" update="true" username="admin"] 1}], New value=["f5226cab-525b-4e33-a26d-
e5053fbba0a1" {"rules":[{"section_id":"f5226cab-525b-4e33-a26d-
e5053fbba0a1","resource_type":"FirewallRule","id":"536870917","display_name":"mp_Rule1_updated
","sources_excluded":false,"destinations_excluded":false,"action":"ALLOW","disabled":false,"lo
gged":false,"direction":"IN_OUT","ip_protocol":"IPV4_IPV6","is_default":false,"_revision":1}],
"resource_type":"FirewallSectionRuleList","id":"f5226cab-525b-4e33-a26d-
e5053fbba0a1","display_name":"FirewallSection-2","section_type":"LAYER3","stateful":true,"rule
_count":1,"is_default":false,"locked":false,"comments":"Default section unlock
comment","lock_modified_by":"admin","lock_modified_time":1597191953299,"autoplumbed":false,"en
forced_on":"VIF","tcp_strict":false,"category":"Default","_protection":"UNKNOWN","_revision":1
}]
```

Deleting a rule (`mp_Rule2`) from a section (`FirewallSection-2`):

```
<182>1 2020-08-12T00:33:58.355Z manager1 NSX 1503 FIREWALL [nsx@6876 audit="true" comp="nsx-
manager" entId="f5226cab-525b-4e33-a26d-e5053fbba0a1" level="INFO"
reqId="2db867e0-0407-44a2-8a6c-96895ff14a2f" splitId="m9SdpPw2" splitIndex="1 of 3"
subcomp="manager" update="true" username="admin"] UserName="admin", ModuleName="Firewall",
Operation="UpdateSectionWithRules", Operation status="success", Old
value=[{"section_id":"f5226cab-525b-4e33-a26d-
e5053fbba0a1","resource_type":"FirewallRule","id":"536870918","display_name":"mp_Rule2","sourc
es_excluded":false,"destinations_excluded":false,"action":"ALLOW","disabled":false,"logged":fa
lse,"direction":"IN_OUT","ip_protocol":"IPV4_IPV6","is_default":false}
{"section_id":"f5226cab-525b-4e33-a26d-
e5053fbba0a1","resource_type":"FirewallRule","id":"536870917","display_name":"mp_Rule1_updated
","sources_excluded":false,"destinations_excluded":false,"action":"ALLOW","disabled":false,"lo
gged":false,"direction":"IN_OUT","ip_protocol":"IPV4_IPV6","is_default":false}
{"locked":false,"comments":"Default section unlock
comment","lock_modified_by":"admin","lock_modified_time":1597191953299,"autoplumbed":false,"en
forced_on":"VIF","tcp_strict":false,"category":"Default","resource_type":"FirewallSection","id
":
<182>1 2020-08-12T00:33:58.355Z manager1 NSX 1503 FIREWALL [nsx@6876 audit="true" comp="nsx-
```

```
manager" entId="f5226cab-525b-4e33-a26d-e5053fbba0a1" level="INFO"
reqId="2db867e0-0407-44a2-8a6c-96895ff14a2f" splitId="m9SdpPw2" splitIndex="2 of 3"
subcomp="manager" update="true" username="admin"] "f5226cab-525b-4e33-a26d-
e5053fbba0a1","display_name":"FirewallSection-2","section_type":"LAYER3","stateful":true,"rule
_count":2,"is_default":false,"_create_user":"admin","_create_time":1597191953297,"_last_modifi
ed_user":"admin","_last_modified_time":1597192378372,"_system_owned":false,"_protection":"NOT_
PROTECTED","_revision":3}], New value=["f5226cab-525b-4e33-a26d-e5053fbba0a1" {"rules":
[{"section_id":"f5226cab-525b-4e33-a26d-
e5053fbba0a1","resource_type":"FirewallRule","id":"536870917","display_name":"mp_Rule1_updated
","sources_excluded":false,"destinations_excluded":false,"action":"ALLOW","disabled":false,"lo
gged":false,"direction":"IN_OUT","ip_protocol":"IPV4_IPV6","is_default":false,"_revision":3}],
"resource_type":"FirewallSectionRuleList","id":"f5226cab-525b-4e33-a26d-
e5053fbba0a1","display_name":"FirewallSection-2","section_type":"LAYER3","stateful":true,"rule
_count":2,"is_default":false,"locked":false,"comments":"Default section unlock
comment","lock_modified_by":
<182>1 2020-08-12T00:33:58.355Z manager1 NSX 1503 FIREWALL [nsx@6876 audit="true" comp="nsx-
manager" entId="f5226cab-525b-4e33-a26d-e5053fbba0a1" level="INFO"
reqId="2db867e0-0407-44a2-8a6c-96895ff14a2f" splitId="m9SdpPw2" splitIndex="3 of 3"
subcomp="manager" update="true" username="admin"]
"admin","lock_modified_time":1597191953299,"autoplumbed":false,"enforced_on":"VIF","tcp_strict
":false,"category":"Default","_protection":"UNKNOWN","_revision":3}]
```

Deleting a section (`FirewallSection-2`) that contains a rule (`mp_Rule1`):

```
<182>1 2020-08-12T00:35:01.304Z manager1 NSX 1503 - [nsx@6876 audit="true" comp="nsx-manager"
level="INFO" reqId="f23e091f-aa6e-47a6-945a-98291cc3f0ba" subcomp="manager" username="admin"]
UserName="admin", ModuleName="NSX-Firewall", Operation="DELETE", Operation status="success",
Old value=[FirewallSectionLock [Id=0ffb0688-9f4e-4096-a19f-2d98ce8cfbeb,
sectionId=f5226cab-525b-4e33-a26d-e5053fbba0a1, sectionRevision=0, locked=false,
comments=Default section unlock comment, created_by=admin, create_time=1597191953299,
last_modified_by=admin, last_modified_time=1597191953299]]
<182>1 2020-08-12T00:35:01.324Z manager1 NSX 1503 FIREWALL [nsx@6876 audit="true" comp="nsx-
manager" entId="f5226cab-525b-4e33-a26d-e5053fbba0a1" level="INFO" reqId="f23e091f-
aa6e-47a6-945a-98291cc3f0ba" subcomp="manager" update="true" username="admin"]
UserName="admin", ModuleName="Firewall", Operation="DeleteSection", Operation
status="success", Old value=[null{"section_id":"f5226cab-525b-4e33-a26d-
e5053fbba0a1","resource_type":"FirewallRule","id":"536870917","display_name":"mp_Rule1_updated
","sources_excluded":false,"destinations_excluded":false,"action":"ALLOW","disabled":false,"lo
gged":false,"direction":"IN_OUT","ip_protocol":"IPV4_IPV6","is_default":false}
{"locked":false,"autoplumbed":false,"enforced_on":"VIF","tcp_strict":false,"category":"Default
","resource_type":"FirewallSection","id":"f5226cab-525b-4e33-a26d-
e5053fbba0a1","display_name":"FirewallSection-2","section_type":"LAYER3","stateful":true,"rule
_count":1,"is_default":false,"_create_user":"admin","_create_time":1597191953297,"_last_modifi
ed_user":"admin","_last_modified_time":1597192438335,"_system_owned":false,"_protection":"NOT_
PROTECTED","_revision":4}], New value=["f5226cab-525b-4e33-a26d-e5053fbba0a1"
{"cascade":true}]
```

Edge firewall changes in Manager mode

Note that log messages for a tier-0 logical router and a tier-1 logical router are similar.

Adding a firewall section (`FirewallSection-1`) for a tier-1 logical router (`myT1_mp`):

```
<182>1 2020-08-12T00:09:55.661Z manager1 NSX 1503 - [nsx@6876 audit="true" comp="nsx-manager"
level="INFO" reqId="14af9252-ddc3-4949-8e01-b2c5676ac258" subcomp="manager" username="admin"]
UserName="admin", ModuleName="NSX-Firewall", Operation="CREATE", Operation status="success",
New value=[FirewallSectionLock [I
d=15b61818-2a65-48cf-a98e-7c2f3fccc845, sectionId=9808d1ec-de08-48b3-8173-12f26fb0ae9c,
sectionRevision=0, locked=false, comments=Default section unlock comment, created_by=admin,
create_time=1597190995659, last_modified_by=admin, last_modified_time=1597190995659]]

<182>1 2020-08-12T00:09:55.687Z manager1 NSX 1503 FIREWALL [nsx@6876 audit="true" comp="nsx-
manager" entId="9808d1ec-de08-48b3-8173-12f26fb0ae9c" level="INFO" reqId="14af9252-
ddc3-4949-8e01-b2c5676ac258" subcomp="manager" update="true" username="admin"]
UserName="admin", ModuleName="Firewall", Operation="AddSection", Operation status="success",
New value=[{"operation":"insert_before","id":"095b443a-115d-4bf7-b4f7-192305321e95"}
{"locked":false,"autoplumbed":false,"tcp_strict":false,"display_name":"FirewallSection-1","app
lied_tos":
[{"target_id":"6562738e-73b9-4f21-9461-460ead581daf","target_display_name":"myT1_mp","target_t
ype":"LogicalRouter","is_valid":true}],"section_type":"LAYER3","stateful":true,"is_default":fa
lse,"_system_owned":false,"_protection":"UNKNOWN","_revision":0}]
```

Adding a rule (`myT1_mp_Rule1`) to a section (`FirewallSection-1`):

```
<182>1 2020-08-12T00:13:44.092Z manager1 NSX 1503 FIREWALL [nsx@6876 audit="true" comp="nsx-
manager" entId="9808d1ec-de08-48b3-8173-12f26fb0ae9c" level="INFO"
reqId="d4e7bdef-0cc6-45e9-8884-061b0f688fec" splitId="snErcGKF" splitIndex="1 of 2"
subcomp="manager" update="true" username="admin"] UserName="admin", ModuleName="Firewall",
Operation="UpdateSectionWithRules", Operation status="success", Old
value=[{"locked":false,"comments":"Default section unlock
comment","lock_modified_by":"admin","lock_modified_time":1597190995659,"autoplumbed":false,"en
forced_on":"LOGICALROUTER","tcp_strict":false,"category":"Default","resource_type":"FirewallSe
ction","id":"9808d1ec-
de08-48b3-8173-12f26fb0ae9c","display_name":"FirewallSection-1","applied_tos":
[{"target_id":"6562738e-73b9-4f21-9461-460ead581daf","target_display_name":"myT1_mp","target_t
ype":"LogicalRouter","is_valid":true}],"section_type":"LAYER3","stateful":true,"rule_count":0,
"is_default":false,"_create_user":"admin","_create_time":1597190995657,"_last_modified_user":"
admin","_last_modified_time":1597190995657,"_system_owned":false,"_protection":"NOT_PROTECTED"
,"_revision":0}], New value=["9808d1ec-de08-48b3-8173-12f26fb0ae9c" {"rules":
[{"display_name":"myT1_mp_Rule1","sources_excluded":false,"destinations_excluded":
```

Updating a rule (from `myT1_mp_Rule1` to `myT1_mp_Rule1_updated`) in a section (`FirewallSection-1`):

```
<182>1 2020-08-12T00:15:31.078Z manager1 NSX 1503 FIREWALL [nsx@6876 audit="true" comp="nsx-
manager" entId="9808d1ec-de08-48b3-8173-12f26fb0ae9c" level="INFO" reqId="eb880eee-5798-42fc-
a8aa-58b70e4aa152" splitId="WviLd4ja" splitIndex="1 of 3" subcomp="manager" update="true"
username="admin"] UserName="admin", ModuleName="Firewall",
Operation="UpdateSectionWithRules", Operation status="success", Old
value=[{"locked":false,"comments":"Default section unlock
comment","lock_modified_by":"admin","lock_modified_time":1597190995659,"autoplumbed":false,"en
forced_on":"LOGICALROUTER","tcp_strict":false,"category":"Default","resource_type":"FirewallSe
ction","id":"9808d1ec-
de08-48b3-8173-12f26fb0ae9c","display_name":"FirewallSection-1","applied_tos":
[{"target_id":"6562738e-73b9-4f21-9461-460ead581daf","target_display_name":"myT1_mp","target_t
```

```
ype":"LogicalRouter","is_valid":true}],"section_type":"LAYER3","stateful":true,"rule_count":1,
"is_default":false,"_create_user":"admin","_create_time":1597190995657,"_last_modified_user":"
admin","_last_modified_time":1597191224058,"_system_owned":false,"_protection":"NOT_PROTECTED"
,"_revision":1}{"section_id":"9808d1ec-
de08-48b3-8173-12f26fb0ae9c","resource_type":"FirewallRule","id":"536870914","display_name":"m
yT1_mp_Rule1","sources_excluded":
<182>1 2020-08-12T00:15:31.078Z manager1 NSX 1503 FIREWALL [nsx@6876 audit="true" comp="nsx-
manager" entId="9808d1ec-de08-48b3-8173-12f26fb0ae9c" level="INFO" reqId="eb880eee-5798-42fc-
a8aa-58b70e4aa152" splitId="WviLd4ja" splitIndex="2 of 3" subcomp="manager" update="true"
username="admin"]
false,"destinations_excluded":false,"action":"ALLOW","disabled":false,"logged":false,"directio
n":"IN_OUT","ip_protocol":"IPV4_IPV6","is_default":false}], New value=["9808d1ec-
de08-48b3-8173-12f26fb0ae9c" {"rules":[{"section_id":"9808d1ec-
de08-48b3-8173-12f26fb0ae9c","resource_type":"FirewallRule","id":"536870914","display_name":"m
yT1_mp_Rule1_updated","sources_excluded":false,"destinations_excluded":false,"action":"ALLOW",
"disabled":false,"logged":false,"direction":"IN_OUT","ip_protocol":"IPV4_IPV6","_revision":1}]
,"resource_type":"FirewallSectionRuleList","id":"9808d1ec-
de08-48b3-8173-12f26fb0ae9c","display_name":"FirewallSection-1","applied_tos":
[{"target_id":"6562738e-73b9-4f21-9461-460ead581daf","target_display_name":"myT1_mp","target_t
ype":"LogicalRouter","is_valid":true}],"section_type":"LAYER3","stateful":true,"is_default":fa
lse,"locked":false,"comments":"Default section unlock
comment","lock_modified_by":"admin","lock_modified_time":1597190995659,"autoplumbed":
<182>1 2020-08-12T00:15:31.078Z manager1 NSX 1503 FIREWALL [nsx@6876 audit="true" comp="nsx-
manager" entId="9808d1ec-de08-48b3-8173-12f26fb0ae9c" level="INFO" reqId="eb880eee-5798-42fc-
a8aa-58b70e4aa152" splitId="WviLd4ja" splitIndex="3 of 3" subcomp="manager" update="true"
username="admin"]
false,"enforced_on":"LOGICALROUTER","tcp_strict":false,"category":"Default","_system_owned":fa
lse,"_protection":"UNKNOWN","_revision":1}]
```


Deleting a rule (`myT1_mp_Rule2`) from a section (`FirewallSection-1`):


```
<182>1 2020-08-12T00:18:05.341Z manager1 NSX 1503 FIREWALL [nsx@6876 audit="true" comp="nsx-
manager" entId="9808d1ec-de08-48b3-8173-12f26fb0ae9c" level="INFO" reqId="bc95016c-5ec2-4b25-
ab17-0b10b6c5a4f0" splitId="damZHQkr" splitIndex="1 of 3" subcomp="manager" update="true"
username="admin"] UserName="admin", ModuleName="Firewall",
Operation="UpdateSectionWithRules", Operation status="success", Old
value=[{"locked":false,"comments":"Default section unlock
comment","lock_modified_by":"admin","lock_modified_time":1597190995659,"autoplumbed":false,"en
forced_on":"LOGICALROUTER","tcp_strict":false,"category":"Default","resource_type":"FirewallSe
ction","id":"9808d1ec-
de08-48b3-8173-12f26fb0ae9c","display_name":"FirewallSection-1","applied_tos":
[{"target_id":"6562738e-73b9-4f21-9461-460ead581daf","target_display_name":"myT1_mp","target_t
ype":"LogicalRouter","is_valid":true}],"section_type":"LAYER3","stateful":true,"rule_count":2,
"is_default":false,"_create_user":"admin","_create_time":1597190995657,"_last_modified_user":"
admin","_last_modified_time":1597191475552,"_system_owned":false,"_protection":"NOT_PROTECTED"
,"_revision":3}{"section_id":"9808d1ec-
de08-48b3-8173-12f26fb0ae9c","resource_type":"FirewallRule","id":"536870914","display_name":
<182>1 2020-08-12T00:18:05.341Z manager1 NSX 1503 FIREWALL [nsx@6876 audit="true" comp="nsx-
manager" entId="9808d1ec-de08-48b3-8173-12f26fb0ae9c" level="INFO" reqId="bc95016c-5ec2-4b25-
ab17-0b10b6c5a4f0" splitId="damZHQkr" splitIndex="2 of 3" subcomp="manager" update="true"
username="admin"]
"myT1_mp_Rule1_updated","sources_excluded":false,"destinations_excluded":false,"action":"ALLOW
","disabled":false,"logged":false,"direction":"IN_OUT","ip_protocol":"IPV4_IPV6","is_default":
false}{"section_id":"9808d1ec-
de08-48b3-8173-12f26fb0ae9c","resource_type":"FirewallRule","id":"536870915","display_name":"m
```

yT1_mp_Rule2","sources_excluded":false,"destinations_excluded":false,"action":"ALLOW","disable
d":false,"logged":false,"direction":"IN_OUT","ip_protocol":"IPV4_IPV6","is_default":false}],
New value=["9808d1ec-de08-48b3-8173-12f26fb0ae9c" {"rules":[{"section_id":"9808d1ec-
de08-48b3-8173-12f26fb0ae9c","resource_type":"FirewallRule","id":"536870914","display_name":"m
yT1_mp_Rule1_updated","sources_excluded":false,"destinations_excluded":false,"action":"ALLOW",
"disabled":false,"logged":false,"direction":"IN_OUT","ip_protocol":"IPV4_IPV6","_revision":3}]
,"resource_type":"FirewallSectionRuleList","id":"9808d1ec-
de08-48b3-8173-12f26fb0ae9c","display_name":"FirewallSection-1","applied_tos"
<182>1 2020-08-12T00:18:05.341Z manager1 NSX 1503 FIREWALL [nsx@6876 audit="true" comp="nsx-
manager" entId="9808d1ec-de08-48b3-8173-12f26fb0ae9c" level="INFO" reqId="bc95016c-5ec2-4b25-
ab17-0b10b6c5a4f0" splitId="damZHQkr" splitIndex="3 of 3" subcomp="manager" update="true"
username="admin"] :
[{"target_id":"6562738e-73b9-4f21-9461-460ead581daf","target_display_name":"myT1_mp","target_t
ype":"LogicalRouter","is_valid":true}],"section_type":"LAYER3","stateful":true,"is_default":fa
lse,"locked":false,"comments":"Default section unlock
comment","lock_modified_by":"admin","lock_modified_time":1597190995659,"autoplumbed":false,"en
forced_on":"LOGICALROUTER","tcp_strict":false,"category":"Default","_system_owned":false,"_pro
tection":"UNKNOWN","_revision":3}]

Deleting a section (`FirewallSection-1`) that contains a rule (`myT1_mp_Rule2`):

<182>1 2020-08-12T00:21:27.646Z manager1 NSX 1503 - [nsx@6876 audit="true" comp="nsx-manager"
level="INFO" reqId="781f43d5-0b4c-494e-89a1-cbc2998fc232" subcomp="manager" username="admin"]
UserName="admin", ModuleName="NSX-Firewall", Operation="DELETE", Operation status="success",
Old value=[FirewallSectionLock [Id=15b61818-2a65-48cf-a98e-7c2f3fccc845, sectionId=9808d1ec-
de08-48b3-8173-12f26fb0ae9c, sectionRevision=0, locked=false, comments=Default section unlock
comment, created_by=admin, create_time=1597190995659, last_modified_by=admin,
last_modified_time=1597190995659]]
<182>1 2020-08-12T00:21:27.669Z manager1 NSX 1503 FIREWALL [nsx@6876 audit="true" comp="nsx-
manager" entId="9808d1ec-de08-48b3-8173-12f26fb0ae9c" level="INFO"
reqId="781f43d5-0b4c-494e-89a1-cbc2998fc232" splitId="u3AofFMr" splitIndex="1 of 2"
subcomp="manager" update="true" username="admin"] UserName="admin", ModuleName="Firewall",
Operation="DeleteSection", Operation status="success", Old
value=[{"locked":false,"autoplumbed":false,"enforced_on":"LOGICALROUTER","tcp_strict":false,"c
ategory":"Default","resource_type":"FirewallSection","id":"9808d1ec-
de08-48b3-8173-12f26fb0ae9c","display_name":"FirewallSection-1","applied_tos":
[{"target_id":"6562738e-73b9-4f21-9461-460ead581daf","target_display_name":"myT1_mp","target_t
ype":"LogicalRouter","is_valid":true}],"section_type":"LAYER3","stateful":true,"rule_count":1,
"is_default":false,"_create_user":"admin","_create_time":1597190995657,"_last_modified_user":"
admin","_last_modified_time":1597191671601,"_system_owned":false,"_protection":"NOT_PROTECTED"
,"_revision":6}{"section_id":"9808d1ec-
de08-48b3-8173-12f26fb0ae9c","resource_type":"FirewallRule","id":"536870916","display_name":"m
yT1_mp_Rule1","sources_excluded":false,"destinations_excluded":false,"action":"ALLOW","disable
d":false,"logged":false,"direction":
<182>1 2020-08-12T00:21:27.669Z manager1 NSX 1503 FIREWALL [nsx@6876 audit="true" comp="nsx-
manager" entId="9808d1ec-de08-48b3-8173-12f26fb0ae9c" level="INFO"
reqId="781f43d5-0b4c-494e-89a1-cbc2998fc232" splitId="u3AofFMr" splitIndex="2 of 2"
subcomp="manager" update="true" username="admin"]
"IN_OUT","ip_protocol":"IPV4_IPV6","is_default":false}null], New value=["9808d1ec-
de08-48b3-8173-12f26fb0ae9c" {"cascade":true}]

Customer Experience Improvement Program

NSX participates in VMware's Customer Experience Improvement Program (CEIP).

Details regarding the data collected through CEIP and the purposes for which it is used by VMware are set forth at the Trust & Assurance Center at https://www.vmware.com/solutions/trustvmware/ceip-products.html.

To join or leave the CEIP for NSX, or edit program settings, see Edit the Customer Experience Improvement Program Configuration.

Edit the Customer Experience Improvement Program Configuration

When you install or upgrade NSX Manager, you can decide to join the CEIP and configure data collection settings.

You can also edit the existing CEIP configuration to join or leave the CEIP program, define the frequency and the days the information is collected, and proxy server configuration.

Prerequisites

- Verify that the NSX Manager is connected and can synchronize with your hypervisor.
- Verify that NSX is connected to a public network for uploading data.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Select **System > General Settings > Customer Program**.
3. Click **Edit** in the Customer Experience Improvement Program section.
4. In the Edit Customer Experience Program dialog box, select the **Join the VMware Customer Experience Improvement Program** check box.
5. Toggle the **Schedule** switch to disable or enable the data collection.

   The schedule is enabled by default.
6. (Optional) Configure the data collection and upload recurrence settings.
7. Click **Save**.

Find the SSH Fingerprint of a Remote Server

Some tasks that involve communication with a remote server require that you provide the SSH fingerprint for the remote server. The SSH fingerprint is derived from a host key on the remote server.

To connect using SSH, the NSX Manager and the remote server must have a host key type in common. Support includes key size 256-bit, 384-bit, and 521-bit. Ensure whatever key size is used at time of backup is used at time of restore. The default location of this key is `/etc/ssh/ssh_host_ecdsa_key.pub`.

Having the fingerprint for a remote server helps you confirm you are connecting to the correct server, protecting you from man-in-the-middle attacks. You can ask the administrator of the remote server to provide the SSH fingerprint of the server. Or you can connect to the remote server to find the fingerprint. Connecting to the server over console is more secure than over the network.

Procedure

1 Log in to the remote server as root.

   Logging in using a console is more secure than over the network.

2 Verify the required hashed ECDSA host key is present on the backup server by running `#ssh-keyscan -t ecdsa <backup server IP/FQDN>`.

```
#ssh-keyscan -t ecdsa ftpserver.corp.local
   #ftpserver.corp.local:22 SSH-2.0-OpenSSH_7.6p1 Ubuntu-4ubuntu0.5
   ftpserver.corp.local ecdsa-sha2-nistp256
```

   If the command output does not return an ECDSA host key, you must configure the key on the backup server. Contact the OS vendor if you need guidance for that configuration.

3 Locate the ECDSA key. The default location of the key is `/etc/ssh/ssh_host_ecdsa_key.pub`.

```
$ ls -al /etc/ssh/*pub
-rw-r--r-- 1 root root  93 Apr  8 18:10 ssh_host_ecdsa_key.pub
-rw-r--r-- 1 root root 393 Apr  8 18:10 ssh_host_rsa_key.pub
```

4 Get the fingerprint of the key.

```
ssh-keygen -lf /etc/ssh/ssh_host_ecdsa_key.pub | awk '{print $2}'
```

Configuring an External Load Balancer

You can configure an external load balancer to distribute traffic to the NSX Managers in a manager cluster.

An NSX Manager cluster does not require an external load balancer. The NSX Manager virtual IP (VIP) provides resiliency in the event of a Manager node failure but has the following limitations:

- VIP does not perform load balancing across the NSX Managers.
- VIP requires all the NSX Managers to be in the same subnet.
- VIP recovery takes about 1 - 3 minutes in the event of a Manager node failure.

An external load balancer can provide the following benefits:

- Load balance across the NSX Managers.
- The NSX Managers can be in different subnets.

- Fast recovery time in the event of a Manager node failure.

An external load balancer will not work with the NSX Manager VIP. Do not configure an NSX Manager VIP if you use an external load balancer.

Authentication Methods When Accessing NSX Manager

The following authentication methods are supported by NSX Manager. For more information about the authentication methods, see the *NSX API Guide*.

- HTTP Basic Authentication
- Session-Based Authentication
- Authentication using an X.509 certificate and a Principal Identity
- Authentication in VMware Cloud on AWS (VMC)

The session-based authentication method (used when you access NSX Manager from a browser) requires source-IP persistence (all requests from the client must go to the same NSX Manager). The other methods do not require source-IP persistence (requests from the client can go to different NSX Managers).

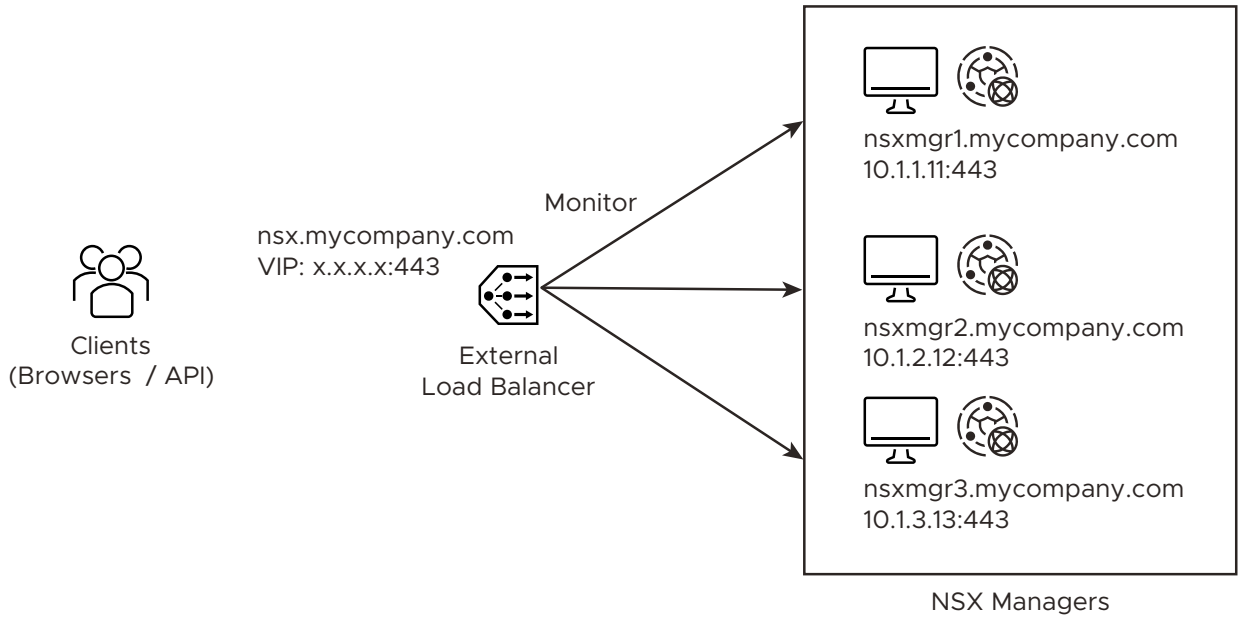
Recommendations

- Create a single VIP on the load balancer with source-IP persistence configured to handle all the authentication methods.
- If you have applications or scripts that might generate a lot of requests to NSX Manager, create a second VIP without source-IP persistence for these applications or scripts. Use the first VIP for browser access to NSX Manager only.

The VIP must have the following configurations:

- Type: Layer4-TCP
- Port: 443
- Pool: NSX Manager Pool
- Persistence: Source-IP persistence for the first VIP. None for the second VIP (if present).

Example of an external load balancer configuration:



NSX Manager's Certificate

The clients access NSX Manager using a FQDN name (for example, nsx.mycompany.com). This FQDN is resolved to the load balancer's VIP. To avoid any certificate mismatch, each NSX Manager must have a certificate that is valid for the VIP's FQDN name. Therefore, you must configure each NSX Manager with a SAN certificate that is valid for its own name (for example, nsxmgr1.mycompany.com) and the VIP's FQDN.

Monitoring the Health of NSX Managers

The load balancer can check that each NSX Manager is running with the following API:

```
GET /api/v1/reverse-proxy/node/health
```

The request headers are:

- Header1
  - Name: Authorization
  - Value: Basic <Base64 Value>

  Note: <Base64 Value> is Username:Password encoded in Base64. You can use https://www.base64encode.net to do the encoding. For example, Header1 might be `Authorization: Basic YWRtaW46Vk13YXJlMSFWTXdhcmUxIQ==` for `admin:VMware1!VMware1!`.
- Header2
  - Name: Content-Type
  - Value: application/json

- Header3
  - Name: Accept
  - Value: application/json

A response indicating that the NSX Manager is running will be:

```
"healthy" : true
```

Note that the format of the response is `"healthy"<space>:<space>true`.

If you change the password of the user that you specify in Header1, you must update Header1 accordingly.

Configure Proxy Settings

You can configure proxy settings for your NSX Manager environment. Do not set up proxy configuration on an NSX Edge, since transport nodes require direct Internet connectivity for FQDN analysis and for URL filtering.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Select **System > General Settings**.
3. To create a proxy configuration the first time, click the highlighted button on the Internet Proxy Server page or continue to the next step.
4. To turn on proxy server configuration, toggle the **Proxy Enabled** to Yes.
5. In the **Scheme** column, select **HTTP** or **HTTPS**.
6. In the **Host** column, enter an IP address.
7. In the **Port** column, enter a port number.
8. In the **Username** column, enter a user name.
9. In the **Password** column, enter a password.
10. Click **Save**.

Promote Manager Objects to Policy Objects

With the business decision to move the consumption layer to policy, the existing configuration needs to be moved from NSX Manager to NSX Policy without data path disruption or deletion or recreation of existing objects. With this feature, you can promote objects created on NSX Manager to NSX Policy and can then later interact with the same objects through NSX Policy UI or NSX Policy APIs.

The promotion process has the following workflow:

1 Collect all manager objects.

2 Translate manager resources to corresponding policy resources intents and apply translated policy resources on policy.

3 Link the obtained policy intents in Step 2 to corresponding existing manager objects.

4 Report policy promotion progress and list the promoted objects.

Promotion of objects occurs based on their dependency order, for example, a group is promoted first and then any rule that consumes that group. Also, note that some configurations and entities are not supported for promotion for any of the following reasons:

- They are policy-only features
- They are not supported on policy yet
- They are deprecated features
- They have passthrough APIs to manager through policy

Objects that are not supported for promotion are as follows:

- AD Configuration
- Policy Based Routing (Forwarding policies)
- L2 forwarder
- LbTcpProfile
- Service insertion
- Traceflow
- End Point protection (Service insertion consumption)
- EVPN and EVPN Tenant
- Gateway QoS Profiles
- Multicast configuration R
- IDS
- Backup restore and proxy settings
- License Management
- Upgrade
- LRQoSProfile
- VRF config on routers
- Bridge Firewall
- Port mirroring session - Local Span and Remote Span

- Multicast config
- OSPF

Supported objects, but unsupported configurations are as follows:

- IP block subnet
- L2 VPN client session

A mixed mode is also not supported for promotion. Mixed mode is where configuration contains combination of policy and manager objects, for example, NAT rules on manager attached to routers created through policy and groups created through policy used in MP DFW rules.

On a Federation setup, you cannot promote objects created on NSX Manager to NSX Policy. If you want to onboard sites to GM in Federation, then first promote all manager objects to policy using this feature. Also, note that for post site and config onboarding this feature is not supported.

When you log in to NSX, an application-level alert is displayed if objects are available for promotion along with a link to initiate the promotion. You can click the link to start the promotion. You can also start the promotion from the **System** tab. If you performed the promotion process earlier, you can also view a history of last five promotions performed and details of data of the last two successful promotions by clicking **Recent Activity**.

Once you initiate the promotion process and the process starts, the system displays a progress bar to show percentage of promotion performed. It also displays manager objects that are promoted to policy objects and status of promotion whether objects succeeded or failed the promotion. You can view failure details by clicking the object failed link against failed objects. Also, if any object fails to get promoted, you can skip it and continue the promotion or you can choose to stop the promotion. If you stop the promotion, the system rollbacks promoted objects to their previous states.

Prerequisites

- You must start the migration coordinator service by running the following command on any one node of manager cluster nodes.

  ```
  start service migration-coordinator
  ```

  Note The entire promotion process will run only on that single node on which you start the migration coordinator service.

- Take a backup before performing the manager to policy promotion. In case a rollback fails, we can revert the system to its original state using the backup.

Procedure

1 Navigate to **System > General Settings > Manager Objects Promotion**

2 Click **Start Objects Promotion**.

  The system displays summary of manager objects.

3 Click **Continue**.

The system starts the promotion and displays the progress and status of promotion. If any object fails to get promoted, the system displays an error. You can click **Skip and Continue** to continue the promotion, or you can click **Cancel** to stop the promotion.

4 Once the promotion is completed successfully, the system displays the Manager to Policy Objects Promotion page.

Back up and restore NSX configured in VMware vCenter

Back up and restore NSX Manager that is configured in VMware vCenter from the vSphere Client.

Note When you restore NSX Manager, use the `nsx-unified-appliance-<releaseversion.buildversion>.ova` file.

Procedure

1 From your browser, log in with admin privileges to a vCenter Server (version 7.0.3 or later) at https://<VMware vCenter>.

2 Click **Login** to **Launch vSphere Client**.

3 In the vSphere Client, from the main menu list, click NSX. The NSX page is displayed.

4 From the vSphere Client, accesss the NSX Manager UI.

5 Configure a backup location and perform backup. See Configure Backups.

6 If the NSX Manager is deleted for some reason, connectivity between the NSX Manager and VMware vCenter is lost.

7 With admin privileges, log in to VMware vCenter, at https://<VMware vCenter>.

8 Install NSX Manager. See *Install NSX Manager and Available Appliances* topic in the *NSX Installation Guide*.

9 Configure a backup location. See Configure Backups.

10 Restore the backup. See Restore a Backup.

Results

Restore is complete. You can access the NSX UI from VMware vCenter.

Configure an External Entropy Source

You can configure an external entropy source to generate random numbers for cryptographic operations. The external entropy source can be used to satisfy one of the requirements of EAL4.

After the ESXi host is configured with an external entropy source, NSX will automatically obtain entropy data from ESXi. No configuration on NSX is required.

The external entropy will impact the functions related to random number generation and cryptographic operation mentioned in Common Criteria Compliance.

Prerequisites

- vSphere 8.0 or later.
- The ESXi host where NSX Manager is running is configured to activate the correct entropy source. For more information, see Controlling ESXi Entropy.

Procedure

1. Log in to vCenter Server from vSphere Client.
2. Browse to the NSX Manager VM.
3. Shut down the VM.
4. Right-click the VM and click **Edit Settings**.
5. Select **VM Options**.
6. Click **Advanced** and click **Edit Configuration**.
7. Set the **isolation.tools.getEntropy.disable** parameter to **`FALSE`**.
8. Click **OK**.

Using NSX Cloud

28

NSX Cloud enables you to manage and secure your public cloud inventory using NSX.

See "Installing NSX Cloud Components" in the *NSX Installation Guide* for the NSX Cloud deployment workflow.

See also: public cloud.

This chapter includes the following topics:

- Cloud Service Manager: UI Walkthrough
- Threat Detection using the NSX Cloud Quarantine Policy
- NSX Enforced Mode
- Native Cloud Enforced Mode
- NSX Features Supported with NSX Cloud
- Deploying NSX Management Components On Microsoft Azure
- Managing Backup and Restore of NSX Manager and CSM in Microsoft Azure
- NSX Cloud FAQs and Troubleshooting

Cloud Service Manager: UI Walkthrough

The Cloud Service Manager (CSM) provides a single pane of glass management endpoint for your public cloud inventory.

The CSM interface is divided into the following categories:

- **Search**: You can use the search text box to find public cloud accounts or related constructs.
- **Clouds**: Your public cloud inventory is managed through the sections under this category.
- **System**: You can access **Settings**, **Utilities**, and **Users** for Cloud Service Manager from this category.

You can perform all public cloud operations by going to the **Clouds** subsection of CSM.

To perform system-based operations, such as, backup, restore, upgrade, and user management, go to the **System** subsection.

Clouds

These are the sections under **Clouds**:

Overview

Access your public cloud account by clicking **Clouds > Overview**.

Each tile on this page represents your public cloud account with the number of accounts, regions, VPCs or VNets, and instances (workload VMs) it contains.

You can perform the following tasks:

Add a public cloud account or subscription	You can add one or more public cloud accounts or subscriptions. This enables you to view your public cloud inventory in CSM. It also shows the number of VMs that are managed by NSX and their state. See "Add Your Public Account" in the *NSX Installation Guide* for instructions.
Deploy/Undeploy NSX Public Cloud Gateway	You can deploy or undeploy one or two (for High Availability) PCG(s). You can also undeploy PCG from CSM. See "Deploy or Link PCGs" in the *NSX Installation Guide* for instructions.
Enable or Disable Quarantine Policy	You can enable or disable Quarantine Policy. See Threat Detection using the NSX Cloud Quarantine Policy for details.
Switch between Grid and Card view	The cards display an overview of your inventory. The grid displays more details. Click the icons to switch between the view types.

CSM provides a single pane of glass view of all your public cloud accounts that you have connected with NSX Cloud by presenting your public cloud inventory in different ways:

- You can view the number of regions you are operating in.
- You can view the number of VPCs/VNets per region.
- You can view the number of workload VMs per VPC/VNet.

There are four tabs under **Clouds**.

Accounts

You can add your public cloud account by navigating to **Clouds > <your public cloud> > Accounts** section of CSM. You can also view information on the public cloud accounts you have already added.

Each card represents a public cloud account of the cloud provider you selected.

You can perform the following actions from this section:

- Add Account
- Edit Account
- Delete Account
- Resync Account

Regions

Navigate to **Clouds > <your public cloud> > Regions** to see your inventory for a selected region.

You can filter the regions by your public cloud account. Each region has VPCs/VNets and instances. If you have deployed any PCGs, you can see them here as **Gateways** with an indicator for the PCG's health.

If you do not have any VPCs/VNets in a public cloud region, that region is not displayed in CSM.

VPCs or VNets

Navigate to **Clouds > <your public cloud> > VPCs** *or* **VNets** to view the VPCs or VNets in your public cloud account or subscription.

You can filter the inventory by account and region.

- Each card represents one VPC/VNet.
- You can have one or two (for HA) PCGs deployed on Transit VPCs/VNets.
- You can link Compute VPCs/VNets to Transit VPCs/VNets.
- You can view more details for each VPC or VNet by switching to the grid view.

  In the grid view you can see three tabs: **Overview**, **Instances**, and **Segments**.

  - **Overview** lists the options under Actions as described in the next step.
  - **Instances** displays a list of instances in the VPC/VNet.
  - **Segments** displays overlay segments in NSX.

    Note This feature is not supported in the current release for NSX Cloud. Do not tag your workload VMs in AWS or Microsoft Azure with tags shown on this screen.

- Click **Actions** to access the following:
  - **Edit Configuration** (only available for Transit VPCs/VNets):
    - Enable or disable Quarantine Policy if in the NSX Enforced Mode.
    - Change your proxy server selection.
  - **Link to Transit VPC/VNet**: This option is only available to VPCs/VNets that do not have any PCG deployed on them. Click to select a Transit VPC/VNet to link to.
  - **Deploy NSX Cloud Gateway**: This option is only available to VPCs/VNets that do not have a PCG deployed on them. Click this option to get started with deploying PCG on this VPC/VNet and make it a Transit or self-managed VPC/VNet. See **Deploy or Link NSX Public Cloud Gateways** in the *NSX Installation Guide* for detailed instructions.

Instances

The **Clouds > <your public cloud> > Instances** section displays details of all instances in your VPC or VNet.

- You can filter the instance inventory by account, region, and VPC or VNet.
- Within your selected account, region, and VPC/VNet, you can further filter instances by the following criteria: **Powered Off**, **Managed**, **Unmanaged**, **Errored**, **User Managed**, **Quarantined**, **Needs Update**, and **Horizon VDI**. Note that Horizon VDI is only available starting NSX 3.1.1 and only for VNets.
- Each card represents an instance (workload VM) and displays a summary.
- For details on the instance, click on the card or switch to the grid view. Among other details, you can see the following:
  - **Rules Realization**: For workload VMs managed in the Native Cloud Enforced Mode, you can see the status of DFW rules created in NSX Manager. The realization status can be *Successful* or *Failed*. You can click on *failed* status to view the error message. See Set up Micro-segmentation for Workload VMs in the Native Cloud Enforced Mode for details.

You can add instances to or remove instances from the CSM User Managed list. See User Managed List for VMs for details.

System

These are the sections under **System**:

System > Settings

These settings are first configured when you install CSM. You can edit them thereafter.

Join CSM with NSX Manager

You must connect the CSM appliance with NSX Manager to allow these components to communicate with each other.

Prerequisites

- NSX Manager must be installed and you must have the username and password for the admin account to log in to NSX Manager.
- CSM must be installed and you must have the Enterprise Administrator role assigned in CSM.

Procedure

1. From a browser, log in to CSM.
2. When prompted in the setup wizard, click **Begin Setup**.

3 Enter the following details in the NSX Manager Credentials screen:

Option	Description
NSX Manager Host Name	Enter the fully qualified domain name (FQDN) of the NSX Manager, if available. You may also enter the IP address of NSX Manager.
Admin Credentials	Enter an Enterprise Administrator username and password for NSX Manager.
Manager Thumbprint	Optionally, enter the NSX Manager's thumbrpint value. If you leave this field blank, the system identifies the thumbprint and displays it in the next screen.

4 (Optional) If you did not provide a thumbprint value for NSX Manager, or if the value was incorrect, the **Verify Thumbprint** screen appears. Select the checkbox to accept the thumbprint discovered by the system.

5 Click **Connect**.

Note If you missed this setting in the setup wizard or if you want to change the associated NSX Manager, log in to CSM, click **System > Settings**, and click **Configure** on the panel titled **Associated NSX Node**.

CSM verifies the NSX Manager thumbprint and establishes connection.

6 (Optional) Set up the Proxy server. See instructions in (Optional) Configure Proxy Servers.

(Optional) Configure Proxy Servers

If you want to route and monitor all internet-bound HTTP/HTTPS traffic through a reliable HTTP Proxy, you can configure up to five proxy servers in CSM.

All public cloud communication from PCG and CSM is routed through the selected proxy server.

Proxy settings for PCG are independent of proxy settings for CSM. You can choose to have none or a different proxy server for PCG.

You can choose the following levels of authentication:

- Credentials-based authentication.
- Certificate-based authentication for HTTPS interception.
- No authentication.

Procedure

1 Click **System > General Settings > Internet Proxy Server**.

Note You can also provide these details when using the CSM Setup Wizard that is available when you first install CSM.

2 On the Internet Proxy Server screen, enter the following details:

Option	Description
Default	Use this radio button to indicate the default proxy server.
Profile Name	Provide a proxy server profile name. This is mandatory.
Proxy Server	Enter the proxy server's IP address. This is mandatory.
Port	Enter the proxy server's port. This is mandatory.
Authentication	Optional. If you want to set up additional authentication, select this check box and provide valid username and password.
Username	This is required if you select the Authentication checkbox.
Password	This is required if you select the Authentication checkbox.
Certificate	Optional. If you want to provide an authentication certificate for HTTPS interception, select this checkbox and copy-paste the certificate in the text box that appears.
No Proxy	Select this option if you do not want to use any of the proxy servers configured.

System > Utilities

The following utilities are available.

Backup and Restore

You can deploy CSM and NSX Manager in your Microsoft subscription. If you have deployed CSM and NSX Manager in your Microsoft subscription, backup and restore is managed using Microsoft Recovery Service Vault. See Managing Backup and Restore of NSX Manager and CSM in Microsoft Azure.

If you have deployed CSM on-prem, see Backup and Restore of the CSM Appliance.

Support Bundle

Click **Download** to retrieve the support bundle for CSM. This is used for troubleshooting. See the *NSX Troubleshooting Guide* for more information.

Backup and Restore of the CSM Appliance

You can back up a CSM appliance and restore from a backup.

Currently only IP address backups and restores are supported for CSM.

Note If you have deployed CSM and NSX Manager in your Microsoft Azure subscription, use the Microsoft Azure Recovery Vault service for CSM's backup and restore. See Managing Backup and Restore of NSX Manager and CSM in Microsoft Azure.

Backing up CSM

You can backup the CSM appliance manually or set up a recurring backup after you configure a backup server.

You can only restore the CSM appliance from an IP address backup. Do not configure FQDN for the CSM appliance.

Prerequisites

- Verify that you have the SSH fingerprint of the backup file server. Only SHA256 hashed ECDSA key is accepted as a fingerprint. See Find the SSH Fingerprint of a Remote Server.
- Ensure that the directory path already exists where you want to store your backups. You cannot use the root directory (/).

Procedure

1. From your browser, log in with admin privileges to CSM at https://<csm-ip-address>.
2. Select **System > Utilities > Tools**.
3. On the **Backup** tab, click **Configure**.
4. Enter the IP address or host name of the backup file server.
5. Change the default port if required.
6. The protocol field is already filled in. Do not change the value.

   SFTP is the only supported protocol.
7. Enter the username and password required to log in to the backup file server.

   The first time you configure a file server, you must provide a password. Subsequently, if you reconfigure the file server, and the server IP (or hostname), port, and user name are the same, you do not need to enter the password again.
8. In the **Destination Directory** field, enter the absolute directory path where the backups will be stored.

   The directory must already exist and cannot be `/`. If the backup file server is a Windows machine, you still use the forward slash when you specify the destination directory. For example, if the backup directory on the Windows machine is `c:\SFTP_Root\backup`, specify `/SFTP_Root/backup` as the destination directory.

   Note The backup process will generate a name for the backup file that can be quite long. On a Windows server, the length of the full path name of the backup file can exceed the limit set by Windows and cause backups to fail. To avoid this issue, see the KB article https://kb.vmware.com/s/article/76528.
9. To encrypt the backups, enter an **Encryption Passphrase**.

   You will need this passphrase to restore a backup. If you forget the passphrase, you cannot restore any backups.
10. Enter the SSH fingerprint of the server that stores the backups.

    You can leave this blank and accept or reject the fingerprint provided by the server.

11 Click the **Schedule** tab.

12 To enable automatic backups, click the **Automatic Backup** toggle.

13 Click **Weekly** and set the days and time of the backup, or click **Interval** and set the interval between backups.

14 Enabling the **Detect NSX configuration change** option will trigger an unscheduled full configuration backup when it detects any runtime or non-configuration related changes, or any change in user configuration.

   You can specify a time interval for detecting database configuration changes. The valid range is 5 minutes to 1,440 minutes (24 hours).

   Note This option can potentially generate a large number of backups. Use it with caution.

15 Click **Save**.

Results

After you configure a backup file server, you can click **Backup Now** to start a backup at any time.

If your backup server is getting full, see instructions for removing backups: Remove Old Backups.

Restoring CSM from a backup

You can restore a CSM appliance if you have a backup.

You must restore a backup on a new installation of CSM. If the old CSM node is still available, you must power it off, before you start the restore process.

Note You can only restore CSM from an IP address backup. FQDN backups are not supported for CSM.

Prerequisites

- Verify that you have the login credential for the backup file server.
- Verify that you have the SSH fingerprint of the backup file server. See Find the SSH Fingerprint of a Remote Server.
- Verify that you have the passphrase of the backup file.

Procedure

1 If the old CSM node is still available, power it off.

2 Deploy a new CSM node with the same IP address of the original CSM node.

3 From your browser, log in with admin privileges to a new CSM appliance.

4 Select **System > Utilities > Tools**.

5 Click the **Restore** tab.

6 Click **Restore Now**. The Restore wizard opens.

7 Select the check box on the **Prerequisites** screen.

8 Provide details of the remote backup server:

   a Enter the IP address or host name.

   b Change the port number, if necessary.

      The default is 22.

   c To log in to the server, enter the user name and password.

   d In the **Backup Directory** text box, enter the absolute directory path where the backups are stored.

   e Enter the passphrase that was used to encrypt the backup data.

   f Enter the SSH fingerprint of the server that stores the backups.

9 Click **Next**.

10 Select a backup. You can also get a list of available backups by logging in to the backup file server. See Listing Available Backups. Replace NSX Manager with CSM in these instructions, for example, when you are asked to log in to the NSX Manager to run a CLI command, log in to CSM instead.

11 Click **Restore**.

   You lose connectivity until the restore completes. The status of the restore operation is displayed. After the restore operation is completed, the Restore Complete screen is displayed, showing the result of the restore, the timestamp of the backup file, and the start and end time of the restore operation. If the restore failed, the screen displays the step where the failure occurred.

   You can also determine the reason for a restore failure by selecting the log files. Run `get log-file syslog` to view the system log file.

   To restart CSM, run the `service nsx-cloud-service-manager restart` command.

   To reboot the CSM node, run the `reboot` command.

12 After the new CSM node is deployed, delete the original CSM VM that you powered down in Step 1.

System > Users

Users are managed using role-based access control (RBAC).

See Chapter 22 Authentication and Authorization for details.

Threat Detection using the NSX Cloud Quarantine Policy

The Quarantine Policy feature in NSX Cloud provides a threat detection mechanism for your NSX-managed workload VMs.

Quarantine Policy is implemented differently in the two VM-management modes.

Table 28-1. Quarantine Policy Implementation in the NSX Enforced Mode and the Native Cloud Enforced Mode

Configurations related to Quarantine Policy	In the NSX Enforced Mode	In the Native Cloud Enforced Mode
Default state	Disabled when deploying PCG using NSX Tools. You can enable it from the PCG-deployment screen or later. See How to Enable or Disable Quarantine Policy.	Always enabled. Cannot be disabled.
Auto-created security groups unique to each mode	All healthy NSX-managed VMs are assigned the `vm-underlay-sg` security group.	`nsx-<NSX GUID>` security groups are created for and applied to NSX-managed workload VMs that are matched with a Distributed Firewall Policy in NSX Manager
Auto-created Public Cloud Security Groups common to both modes:	The **gw** security groups are applied to the respective PCG interfaces in AWS and Microsoft Azure. ■ gw-mgmt-sg ■ gw-uplink-sg ■ gw-vtep-sg The **vm** security groups are applied to NSX-managed VMs depending on their current state and whether Quarantine Policy is enabled or disabled: ■ `default-vnet-<vnet-id>-sg` in Microsoft Azure and `default` in AWS. **Note** In AWS, the `default` security group already exists. It is not created by NSX Cloud.	

General Recommendation for NSX Enforced Mode :

Start with *disabled* for **Brownfield** deployments: Quarantine Policy is disabled by default. When you already have VMs set up in your public cloud environment, use the disabled mode for Quarantine Policy until you onboard your workload VMs. This ensures that your existing VMs are not automatically quarantined.

Start with *enabled* for **Greenfield** deployments: For greenfield deployments, it is recommended that you enable Quarantine Policy to allow threat detection for your VMs to be managed by NSX Cloud.

Quarantine Policy in the NSX Enforced Mode

Enabling Quarantine Policy is optional in the NSX Enforced Mode.

How to Enable or Disable Quarantine Policy

In the NSX Enforced Mode, you can elect to enable Quarantine Policy in two ways.

The first possibility to enable Quarantine Policy is when you deploy PCG on a Transit VPC/VNet or link a Compute VPC/VNet to a Transit. Move the slider for **Quarantine Policy on the Associated VPC/VNet** to **Enabled** from the default **Disabled** state. See **Deploy PCG** in the *NSX Installation Guide*.

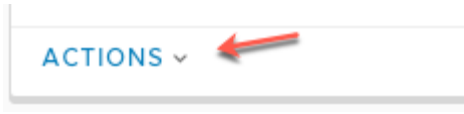
You can also enable Quarantine Policy later following the steps here.

Prerequisites

If enabling Quarantine Policy after deploying or linking to a PCG, you must have one or more Transit or Compute VPCs/VNets onboarded in the NSX Enforced Mode, that is you elected to use NSX Tools for managing your workload VMs.

Procedure

1. Log in to CSM and go to your public cloud:
   a. If using AWS, go to **Clouds > AWS > VPCs**. Click on the Transit or Compute VPC.
   b. If using Microsoft Azure, go to **Clouds > Azure > VNets**. Click on the Transit or Compute VNet.
2. Enable the option using any one of the following:
   - In the tile view, click on **ACTIONS > Edit Configuration**

   - If you are in the grid view, select the checkbox next to the VPC or VNet and click **ACTIONS > Edit Configuration** 

   - If you are in the VPC or VNet's page, click the ACTIONS icon to go to **Edit Configurations**

     
3. Turn **Default Quarantine** on or off.
4. Click **SAVE**.

Quarantine Policy Impact when Disabled

NSX Cloud does not manage the public cloud security groups of untagged VMs when Quarantine Policy is disabled.

However, for VMs tagged with `nsx.network=default` in the public cloud, NSX Cloud assigns appropriate security groups depending on the VM's state. This behavior is similar to when the Quarantine Policy is enabled, but the rules in the quarantine security groups: `default-vnet-<vnet-id>-sg` in Microsoft Azure and `default` in AWS are configured similar to default public cloud security groups, allowing everything within the VPC/VNet and denying all other inbound traffic. Any manual changes to the security groups of tagged VMs are reverted to the NSX Cloud-assigned security group within two minutes.

Note If you do not want NSX Cloud to assign security groups to your NSX-managed (tagged) VMs, add them to the User Manged list in CSM. See User Managed List for VMs.

The following table shows how NSX Cloud manages the public cloud security groups of workload VMs when Quarantine Policy is disabled.

Table 28-2. NSX Cloud assignment of public cloud security groups when Quarantine Policy is disabled

Is VM tagged with *nsx.network=default* in the public cloud?	Is VM added to the User Managed List?	VM's Public cloud security group when Quarantine Policy is disabled and explanation
VM could be tagged or not tagged	Added to the User Managed list.	Retains existing public cloud security group because NSX Cloud doesn't take any action on VMs in the User Managed list.
Not tagged	Not added to the User Managed List	Retains existing public cloud security group because NSX Cloud doesn't take action on untagged VMs.
Tagged	Not added to the User Managed List	■ If VM has no threats: `vm-underlay-sg` ■ If VM has potential threats (see note): `default-vnet-<vnet-id>-sg` in Microsoft Azure; `default` in AWS **Note** The assignment of public cloud security groups is triggered within 90 seconds of applying the `nsx.network=default` tag to your workload VMs. You still need to install NSX Tools for the VMs to be NSX-managed. Until NSX Tools are installed, your tagged workload VMs remain in the default security group. ,

The following table shows how NSX Cloud manages the public cloud security groups of VMs if Quarantine policy was enabled before and is now disabled:

Table 28-3. NSX Cloud assignment of public cloud security groups when Quarantine Policy is disabled from being enabled at first

Is VM tagged with *nsx.network=default* in the public cloud?	**Is VM in the User Managed list?**	**VM's existing public cloud security group when Quarantine Policy is enabled**	**VM's public cloud security group after Quarantine Policy is disabled**
VM could be tagged or not tagged	Yes, VM is in the User Managed list	Any existing public cloud security group	Retains existing public cloud security group because NSX Cloud doesn't take any action on VMs in the User Managed list. **Note** If you have a VM in the User Managed list in any NSX Cloud-assigned security groups, you must manually move it to `default` security group in AWS and `default-vnet-<vnet-id>-sg` security group in Microsoft Azure.
Not tagged	Not added to the User Managed List	`default-vnet-<vnet-id>-sg` (Microsoft Azure) Or `default`(AWS)	Remains in the existing security groups when disabling the Quarantine Policy because it is untagged and not considered NSX-managed. You can manually assign any other security group to this VM as required.
Tagged	Not added to the User Managed List	`vm-underlay-sg` Or `default-vnet-<vnet-id>-sg` (Microsoft Azure) Or `default`(AWS)	Retains the NSX Cloud-assigned security group because that is consistent for tagged VMs in the Quarantine enabled or disabled modes.

Quarantine Policy Impact when Enabled

NSX Cloud manages the public cloud security group of all workload VMs in this VPC/VNet when Quarantine Policy is enabled.

Any manual changes to the security groups are reverted to the NSX Cloud-assigned security group within two minutes. If you do not want NSX Cloud to assign security groups to your VMs, add them to the User Managed list in CSM. See User Managed List for VMs.

Note Removing the VM from the User Managed list causes the VM to revert to the NSX Cloud-assigned security group.

Table 28-4. NSX Cloud assignment of public cloud security groups when Quarantine Policy is enabled

Is VM tagged with *nsx.network=default* in the public cloud)?	**Is VM in the User Managed list?**	**VM's public cloud security group when Quarntine Policy is enabled and explanation**
Tagged	Not added to the User Managed List	■ If VM has no threats: `vm-underlay-sg` ■ If VM has potential threats (see note): `default-vnet-<vnet-ID>-sg` in Microsoft Azure; `default` in AWS **Note** The assignment of public cloud security groups is triggered within 90 seconds of applying the `nsx.network=default` tag to your workload VMs. You still need to install NSX Tools for the VMs to be NSX-managed. Until NSX Tools are installed your tagged workload VMs are quarantined. ,
Not Tagged	Not added to the User Managed List	`default-vnet-<vnet-ID>-sg` in Microsoft Azure; `default` in AWS. Untagged VMs are considered unmanaged and therefore quarantined by NSX Cloud.
Tagged	Yes, VM is in the User Managed List	Retains existing public cloud security group because NSX Cloud doesn't take action on VMs in the User Managed list.
Not Tagged		

The following table captures the impact on security group assignments if the Quarantine Policy was disabled at first and then you enable it:

Table 28-5. NSX Cloud assignment of public cloud security groups when Quarantine Policy is enabled from being disabled at first

Is VM tagged with *nsx.network=default* in the public cloud?	**Is VM in the User Managed list?**	**VM's existing public cloud security group when Quarantine Policy is disabled**	**VM's public cloud security group after Quarantine Policy is enabled**
Not Tagged	Not added to the User Managed List	Any existing public cloud security group	`default-vnet-<vnet-ID>-sg` (Microsoft Azure) Or `default`(AWS)
Tagged	Not added to the User Managed List	`vm-underlay-sg` Or `default-vnet-<vnet-ID>-sg` (Microsoft Azure) Or `default`(AWS)	Retains the NSX Cloud-assigned security group that is consistent for tagged VMs in the Quarantine enabled or disabled modes.

Table 28-5. NSX Cloud assignment of public cloud security groups when Quarantine Policy is enabled from being disabled at first (continued)

Is VM tagged with *nsx.network=default* in the public cloud?	Is VM in the User Managed list?	VM's existing public cloud security group when Quarantine Policy is disabled	VM's public cloud security group after Quarantine Policy is enabled
Tagged	Yes, VM is in the User Managed List	Any existing public cloud security group.	Retains existing public cloud security group because NSX Cloud doesn't take any action on VMs in the User Managed list.
Not Tagged			

Quarantine Policy in the Native Cloud Enforced Mode

Quarantine Policy is always enabled in the Native Cloud Enforced Mode.

Table 28-6. Assignment of Public Cloud Security Groups in the Native Cloud Enforced Mode

Is VM part of a valid NSX Security policy?	Is VM added to the User Managed List?	VM's public cloud security group and explanation
Yes, VM is matched with a valid NSX Security Policy	Not added to User Managed List	NSX Cloud-created public cloud security group named like `nsx-{NSX-GUID}` which is the corresponding public cloud security group for the NSX Security Policy.
No, VM does not have a valid NSX firewall policy	Not added to User Managed List	`default-vnet-<vnet-ID>-sg` in Microsoft Azure or `default` in AWS because this is the threat detection behavior of NSX Cloud. In the Native Cloud Enforced Mode, the NSX Cloud-created security groups `default-vnet-<vnet-ID>-sg` in Microsoft Azure or `default` in AWS mimic the default public cloud security policy. **Note** In CSM the VM shows an Error state.
Yes, VM has valid NSX Security policy	Added to User Managed list	Retains existing public cloud security group because NSX Cloud doesn't take any action on VMs added to the User Managed list.
No, VM does not have a valid NSX Security policy		

User Managed List for VMs

Adding VMs to the **User Managed** list is an option available from CSM for all workload VMs in your public cloud inventory.

You can add VMs to the **User Managed** in both the VM-management modes: NSX Enforced Mode and the Native Cloud Enforced Mode.

Why to add VMs to the User Managed list?

- In the NSX Enforced Mode: If you have the Quarantine Policy enabled and you need to verify any specific DFW policies with existing applications on the VM, add such a VM to the **User Managed** list before onboarding it with NSX Cloud.

- In either the NSX Enforced Mode or the Native Cloud Enforced Mode:
  - If you have VMs with errors and want to access them to resolve the errors, add such VMs to the **User Managed** list so you can move them out of the quarantine state and use debugging tools as required.
  - Add VMs to the **User Managed** list in your public cloud inventory that you don't want NSX to manage, e.g. DNS Forwarder, Proxy server etc.

How to use the User Managed List

Follow these instructions to add VMs to the **User Managed** list or remove them.

Prerequisites

You must have one or more public cloud accounts added to CSM.

Procedure

1. Log in to CSM using an Enterprise Admin account and go to your public cloud account.
   a. If using AWS, go to **Clouds > AWS > VPCs > Instances**.
   b. If using Microsoft Azure, go to **Clouds > Azure > VNets > Instances**.
2. If in Tiles mode, switch to Grid mode by clicking the mode selector in the right corner of the instances view.
3. Select the VMs (instances) that you want to add to or remove from the **User Managed** list.
4. Click **Actions** and select either **Add to User Managed List** or **Remove from User Managed List**.
5. Go back to the Accounts tab, select the account tile and click **Actions > Resync Account**.

Results

Each VM added to the **User Managed** list remains in the security group it was assigned before adding to the **User Managed** list. You can now apply any other security group to the VM as required. NSX Cloud ignores VMs in the **User Managed** list regardless of the status of Quarantine Policy.

If you remove a VM from the **User Managed** list in the Native Cloud Enforced Mode or remove an NSX-managed VM from the **User Managed** list in the NSX Enforced Mode, NSX Cloud starts assigning security groups to that VM depending on its state.

NSX Enforced Mode

In the NSX Enforced Mode, that is, by using NSX Tools, you must first onboard VMs by tagging them in the public cloud and installing NSX Tools on them, before starting to manage these VMs using NSX.

Supported Operating Systems for Workload VMs

This is the list of operating systems currently supported by NSX Cloud for your workload VMs in the NSX Enforced Mode.

See the NSX Cloud Known Issues section in the *NSX Release Notes* for exceptions. For supported operating systems it is assumed that you are using the standard Linux kernel versions. Public cloud marketplace images with custom kernels, for example, upstream Linux kernel with modified sources, are not supported.

On Red Hat Enterprise Linux and its derivatives, SELinux is not supported. To install NSX Tools, disable SELinux.

RHEL Extended Update Support (EUS) kernel in RHEL and CentOS are not supported.

OS	Kernal/Build Versions	Supported In
CentOS 7.2		2.4 and later
CentOS 7.3		2.4 and later
CentOS 7.4		2.4 and later
CentOS 7.5		2.4 and later
CentOS 7.6		2.5 and later
CentOS 7.7		3.2 and later
CentOS 7.8		3.2 and later
CentOS 7.9		3.2 and later
CentOS 8.0		3.2 and later
CentOS 8.1		3.2 and later
CentOS 8.2		3.2 and later
CentOS 8.3		3.2 and later
RHEL 7.0	3.10.0-123	3.1.2 only New deployments are not supported for this OS. Updates may not be available in the future releases.
RHEL 7.1	3.10.0-229	3.1.2 to 3.1.3 New deployments are not supported for this OS. Updates may not be available in the future releases.
RHEL 7.2	3.10.0-327	2.4 and later

OS	Kernal/Build Versions	Supported In
RHEL 7.3	3.10.0-514	2.4 and later
RHEL 7.4	3.10.0-693	2.4 and later
RHEL 7.5	3.10.0-862	2.4 and later
RHEL 7.6	3.10.0-957	2.5 and later
RHEL 7.7	3.10.0-1062	3.0 and later
RHEL 7.8	3.10.0-1127	3.1.2 and later
RHEL 7.9	3.10.0-1160	3.1.2 and later
RHEL 8.0	4.18.0-80	3.1.2 and later
RHEL 8.1	4.18.0-147	3.1.2 and later
RHEL 8.2	4.18.0-193	3.1.2 and later
RHEL 8.3	4.18.0-240	3.1.2 and later
SUSE Linux Enterprise Server (SLES) 12 SP3		3.0 to 3.1 New deployments are not supported for this OS. Updates may not be available in the future releases.
SLES 12 SP4		3.2 and later
Ubuntu 14.04		2.4 to 3.1 New deployments are not supported for this OS. Updates may not be available in the future releases.
Ubuntu 16.04		2.4 to 3.2 New deployments are not recommended for this OS. Updates may not be available in the future releases.
Ubuntu 18.04		2.5 and later
Ubuntu 20.04		3.2 and later
Windows Server 2012 R2 Datacenter	Version 6.3	2.4 and later
Windows Server 2012 R2 Standard	Version 6.3	3.1.2 and later
Windows Server 2016 Datacenter	Version 1607, 1803 and 1809	2.4 and later
Windows Server 2016 Standard	Version 1607	3.1.2 and later

OS	Kernal/Build Versions	Supported In
Windows Server 2019 Datacenter	Version 1809	2.5 and later
Windows Server 2019 Standard	Version 1809	3.1.2 and later
Windows 10 Enterprise	Version 1607, 1803, 1809, 1909, 2004, and 20H2	3.2.0 and later
Windows 10 Home	Version 1607, 1803, 1809, 1909, 2004, and 20H2	3.2.0 and later
Windows 10 Pro	Version 1607, 1803, 1809, 1909, 2004, and 20H2	3.2.0 and later

Note In NSX 3.1.1, we support Red Hat Enterprise Linux (RHEL) 7.8 and 7.9, with an edit required in the NSX Tools installation script.

Replace the following line:

```
if [ "$_ver" != "7.2" -a "$_ver" != "7.3" -a "$_ver" != "7.4" -a "$_ver" != "7.5" -a "$_ver" !
= "7.6" -a "$_ver" != "7.7" ]; then
```

...with this update for RHEL 7.8 and RHEL 7.9:

```
if [ "$_ver" != "7.2" -a "$_ver" != "7.3" -a "$_ver" != "7.4" -a "$_ver" != "7.5" -a "$_ver" !
= "7.6" -a "$_ver" != "7.7" -a "$_ver" != "7.8" -a "$_ver" != "7.9" ]; then
```

Onboarding VMs in the NSX Enforced Mode

Refer to this workflow for an overview of the steps involved in onboarding and managing workload VMs from your public cloud in the NSX Enforced Mode.

Table 28-7. Day-N Workflow for onboarding your workload VMs into NSX Cloud

Task	Instructions
❑ Tag workload VMs with the key-value **`nsx.network=default`**.	Follow instructions in your public cloud documentation for tagging workload VMs.
❑ Install NSX Tools on your Windows and Linux workload VMs. **Note** If **Auto- Install NSX Tools** is enabled in CSM for Microsoft Azure VNets, NSX Tools are automatically installed.	See Install NSX Tools
❑ (Optional) If you have added VMs to the User Managed list in CSM, remove those VMs from the User Managed list that you want to bring under NSX management.	See How to use the User Managed List.

Tag VMs in the Public Cloud

Apply the **nsx.network=default** tag to VMs that you want to manage using NSX.

Procedure

1. Log in to your public cloud account and go to your VPC or VNet where you want your workload VMs to be managed by NSX.
2. Select the VMs that you want to manage using NSX.
3. Add the following tag details for the VMs and save your changes.

```
Key: nsx.network
Value: default
```

Note Apply this tag at the VM level.

Results

You may have already onboarded the VPCs/VNets where you applied the `nsx.network=default` tags to workload VMs. You can also onboard these VPCs/VNets after applying the tag. Successful onboarding of the VPC/VNet results in the workload VMs to be considered NSX-managed.

What to do next

Install NSX Tools on these VMs. See Install NSX Tools.

If using Microsoft Azure, you have the option to auto-install NSX Tools on tagged VMs. See Install NSX Tools Automatically for details.

Install NSX Tools

Install NSX Tools on your workload VMs

There are several options available to install NSX Tools:

- Download and install NSX Tools in individual workload VMs. Linux and Windows VMs have some variations.
- Use replicable images with NSX Tools installed on them using your public cloud's supported method, for example, create an AMI in AWS or a Managed Image in Microsoft Azure.
- AWS-only: When launching VMs, provide the NSX Tools download location and installation command in **User Data**.

- Microsoft Azure-only: Enable auto-installation of NSX Tools when deploying PCG in a Microsoft Azure VNet or while linking to a Transit VNet, or by editing a Transit/Compute VNet's Configuration.

Note If workload VMs on which you want to install NSX Tools are in the User Managed list, ensure the following ports are open in the security groups you have assigned to such VMs:

- Inbound UDP 6081 : For overlay data packets. This should be allowed for (Active/Standby) PCG's VTEP IP address (eth1 interface).
- Outbound TCP 5555 : For control packets. This should be allowed for (Active/Standby) PCG's management IP address (eth0 interface).
- TCP 8080 : For install/upgrade on the PCG's management IP address.
- TCP 80: For downloading any third party dependencies while installing NSX Tools.
- UDP 67,68: For DHCP packets.
- UDP 53: For DNS resolution.

Install NSX Tools on Linux VMs

To install NSX Tools on your Linux workload VMs, follow these instructions.

See Supported Operating Systems for Workload VMs for a list of Linux distributions currently supported.

Note To verify the checksum of this script, go to **VMware Downloads > Drivers & Tools > NSX Cloud Scripts**.

Prerequisites

- Ensure that the following commands are available on your workload VM to run the NSX Tools installation script:
  - **wget**
  - **nslookup**
  - **dmidecode**
- In Microsoft Azure, if you have a custom DNS server configured for the Azure VNet where the workload VMs are running, do one of the following:
  - 1 Add a DNS record to resolve the PCG's private IP and use that DNS name for the download/installation commands.
    2 Provide the corresponding DNS suffix and DNS server parameters for the NSX Tools installation command.
  - Or set up DNS forwarding on your DNS server to forward DNS queries to the PCG's FQDN (nsx-gw.vmware.local) to the Azure DNS: 168.63.129.16

Procedure

1 Log in to CSM and go to your public cloud:

   a If using AWS, go to **Clouds > AWS > VPCs**. Click a Transit or Compute VPC.

   b If using Microsoft Azure, go to **Clouds > Azure > VNets**. Click the VNet on which one or a pair of PCGs is deployed and running.

   Note: Transit VPC/VNet is where one or a pair of PCGs is deployed and running. Compute VPC/VNet is the one linked to a Transit and can use the PCG instances deployed there.

2 From the **NSX Tools Download & Installation** section of the screen, make a note of the **Download Location** and the **Installation Command** under **Linux**.

   Note For VNets, if you have a custom DNS configured, see prerequisites regarding DNS settings.

3 Log in to the Linux workload VM with superuser privileges.

4 Use `wget` or equivalent to download the installation script on your Linux VM from the **Download Location** you noted from CSM. The installation script is downloaded in the directory where you run the `wget` command.

   Note To verify the checksum of this script, go to **VMware Downloads > Drivers & Tools > NSX Cloud Scripts**.

5 Change permissions on the installation script to make it executable if necessary, and run it:

```
$ chmod +x install_nsx_vm_agent.sh && sudo ./install_nsx_vm_agent.sh
```

   Note: On Red Hat Enterprise Linux and its derivatives, SELinux is not supported. To install NSX Tools, disable SELinux.

6 You lose connectivity with your Linux VM after installation of NSX Tools begins. Messages such as the following appear on your screen: `Installation completed!!! Starting NSX Agent service. SSH connection will now be lost.`. To complete the onboarding process, log in to your VM again.

Results

NSX Tools are installed on your workload VM.

Note

- After NSX Tools are successfully installed, port 8888 shows as open on the workload VM but it is blocked for VMs in the underlay mode and must be used only when required for advanced troubleshooting. You can access workload VMs over port 8888 using a jumphost if the jumphost is also in the same VPC as the workload VMs that you want to access.
- The script uses `eth0` as the default interface.

What to do next

Managing VMs in the NSX Enforced Mode

Install NSX Tools on Windows VMs

Follow these instructions to install NSX Tools on your Windows workload VM.

See Supported Operating Systems for Workload VMs for a list of Microsoft Windows versions currently supported.

Note To verify the checksum of this script, go to **VMware Downloads > Drivers & Tools > NSX Cloud Scripts**.

Prerequisites

- In Microsoft Azure, if you have a custom DNS server configured for the Azure VNet where the workload VMs are running, do one of the following:
  - 1 Add a DNS record to resolve the PCG's private IP and use that DNS name for the download/installation commands.
    2 Provide the corresponding DNS suffix and DNS server parameters for the NSX Tools installation command.
  - Or set up DNS forwarding on your DNS server to forward DNS queries to the PCG's FQDN (nsx-gw.vmware.local) to the Azure DNS: 168.63.129.16

Procedure

1 Log in to CSM and go to your public cloud:
   a If using AWS, go to **Clouds > AWS > VPCs**. Click on a Transit or Compute VPC.
   b If using Microsoft Azure, go to **Clouds > Azure > VNets**. Click on the VNet on which one or a pair of PCGs is deployed and running.

   Note: Transit VPC/VNet is where one or a pair of PCGs is deployed and running. Compute VPC/VNet is the one linked to a Transit and can use the PCGs deployed there.

2 From the **NSX Tools Download & Installation** section of the screen, make a note of the **Download Location** and the **Installation Command** under **Windows**.

   Note For VNets, if you have a custom DNS configured, see prerequisite regarding DNS settings.

3 Connect to your Windows workload VM as Administrator.

4 Download the installation script on your Windows VM from the **Download Location** you noted from CSM. You can use any browser, for example, Internet Explorer, to download the script. It is downloaded in your browser's default downloads directory, for example, *C:\Downloads*.

Note To verify the checksum of this script, go to **VMware Downloads > Drivers & Tools > NSX Cloud Scripts**

Note:

5 Open a PowerShell prompt and go to the directory containing the downloaded script.

6 Use the **Installation command** you noted from CSM to run the downloaded script.

For example:

```
c:\> powershell -file 'nsx_install.ps1" -operation install -dnsSuffix <>
```

Note The file argument needs the full path unless you are in the same directory or if the PowerShell script is already in the path. For example, if you download the script to *C:\Downloads*, and you are currently not in that directory, then the script must contain the location: *powershell -file 'C:\Downloads\nsx_install.ps1' ...*

7 The script runs and when completed, displays a message indicating whether NSX Tools was installed successfully.

Note The script considers the primary network interface as the default.

What to do next

Managing VMs in the NSX Enforced Mode

Generate Replicable Images

You can generate an AMI in AWS or a Managed Image in Microsoft Azure of a VM with the NSX agent installed on it.

With this feature, you can launch multiple VMs with with the agent configured and running.

There are two ways in which you can generate an AMI/Managed Image (image in the rest of this topic) of a VM with the NSX agent installed on it:

- **Generate image with an unconfigured NSX agent**: You can generate an image from a VM that has the NSX agent installed on it but not configured by using the `-noStart` option. This option allows the NSX agent package to be fetched and installed but the NSX services are not started. Also, no NSX configurations such as certificate generation, are made.
- **Generate image after removing existing NSX agent configurations**: You can remove configurations from an existing NSX-managed VM and use it for generating an image.

Generating AMI with an unconfigured NSX agent

You can generate an AMI of a VM with the NSX agent installed on it and not configured.

To generate an image from a VM that has the NSX agent installed on it using the **-noStart** option, do the following:

Procedure

1. Copy paste the NSX agent Installation Command from CSM. See instructions at Install NSX Tools
   a. Edit the command for Windows as follows:

   ```
   c:\> powershell -file 'nsx_install.ps1" -operation install -dnsSuffix <> -noStart true
   ```

   b. Edit the command for Linux as follows:

   ```
   $ chmod +x install_nsx_vm_agent.sh && sudo ./install_nsx_vm_agent.sh --no-start
   ```

2. Go to this VM in your public cloud and create an image.

Generating an Image After Removing Existing NSX Agent Configurations

You can generate an image of a VM that has a configured NSX agent.

To remove configurations from an existing NSX-managed VM and use it for generating images, do the following:

Procedure

1. Removing NSX agent configurations from a Windows or Linux VM:
   a. Log in to the workload VM using preferably using a jumphost.
   b. Open the NSX CLI:

   ```
   sudo nsxcli
   ```

   c. Enter the following commands:

   ```
   hostname> set debug
   hostname> clear nsx-vm-agent state
   ```

2. Locate this VM in your public cloud and create an image.

Install NSX Tools Automatically

Currently only supported for Microsoft Azure.

In Microsoft Azure, if the following criteria are met, NSX Tools are installed automatically:

- Azure VM Extensions are installed on the VMs in the VNet added into NSX Cloud. See Microsoft Azure documentation on VM Extensions for more details.

- The security group applied to VMs in Microsoft Azure must allow access to install NSX Tools. If Quarantine Policy is enabled, you can add your VMs to the User Managed list in CSM before installation and remove them from the User Managed list after installation.
- VMs tagged using the key `nsx.network` and value `default`.

To enable this feature:

1 Go to **Clouds > Azure > VNets**.

2 Select the VNet on whose VMs you want to auto-install CSM.

3 Enable the option using any one of the following:

- In the tile view, click on **ACTIONS > Edit Configuration**



- If you are in the grid view, select the checkbox next to the VNet and click **ACTIONS > Edit Configuration**. 

- If you are in the VNet tab, click the ACTIONS icon to go to **Edit Configurations**  .

4 Move the slider next to **Auto-Install NSX Tools** to the ON position.

Note If NSX Tools installation fails, do the following:

1 Log in to the Microsoft Azure portal and navigate to the VM where NSX Tools installation failed.

2 Go to the VM's Extensions and uninstall the extension named `VMwareNsxAgentInstallCustomScriptExtension`.

3 Remove the `nsx.network=default` tag from this VM.

4 Add the `nsx.network=default` tag on this VM again.

Within about three minutes, NSX Tools are installed on this VM.

Install NSX Tools with User Data in AWS

When launching a new workload VM in an AWS VPC, you can install NSX Tools by providing the NSX Tools download and installation instructions in the User Data field.

When you launch an AWS EC2 instance, you have the option of passing `user data` to the instance that can be used to perform common automated configuration tasks, including running scripts after the instance starts. You can pass two types of user data to AWS EC2: shell scripts and cloud-init directives.

Copy the download and installation instructions for NSX Tools from CSM and paste into User Data when launching a new workload VM.

Prerequisites

Before installing NSX Tools using `User Data`, ensure that the Transit and Compute VPCs are peered. This is required so that FQDN given in the download command, such as, `nsx-gw.vmware.local` can be resolved from the launching instance.

Procedure

1 Log in to AWS console and start the process of launching a new workload VM.

2 In another browser window, log in to CSM.

   a Go to **Clouds > AWS > VPCs**

   Note Transit VPC/VNet is where one or a pair of PCGs is deployed and running. Compute VPC/VNet is the one linked to a Transit and can use the PCGs deployed there.

   b Click on a Transit or Compute VPC.

   c From the **NSX Tools Download & Installation** section of the screen, copy the **Download Location** and the **Installation Command** under **Linux** or **Windows** depending on what OS you are using for your workload VM. You can also copy-paste the following shell script:

```
#!/bin/bash
sudo wget http://nsx-gw.vmware.local:8080/factory_default/linux/install_nsx_vm_agent.sh
sudo chmod +x install_nsx_vm_agent.sh && sudo ./install_nsx_vm_agent.sh
```

3 In AWS, in the steps for launching a new workload VM instance, paste the download location and the installation command as **Text** in User Data in the Advanced Details section.

Results

The workload VM is launched and NSX Tools are installed on it automatically.

Uninstalling NSX Tools

Use these OS-specific commands to uninstall NSX Tools.

Uninstalling NSX Tools from a Windows VM

Note To see other options available for the installation script, use `-help`.

1 Remote log in to the VM using RDP.

2 Run the installation script with the uninstall option:

```
\nsx_install.ps1  -operation uninstall
```

Uninstalling NSX Tools from a Linux VM

Note To see other options available for the installation script, use `--help`.

1 Remote log in to the VM using SSH.

2 Run the installation script with the uninstall option:

```
sudo ./install_nsx_vm_agent.sh --uninstall
```

Security Groups after Onboarding in the NSX Enforced Mode

The following security group configurations take place automatically:

If Quarantine Policy is enabled:

- Healthy NSX-managed VMs are moved to the `vm-underlay-sg` in the public cloud.
- Unmanaged VMs or NSX-managed VMs with errors are moved to the `default` Security Group in AWS and `default-vnet-<vnet-ID>-sg` Network Security Group in Microsoft Azure.
- VMs in the User Managed list are not affected.

If Quarantine Policy is disabled:

- Healthy NSX-managed VMs are moved to the `vm-underlay-sg` in the public cloud.
- NSX-managed VMs with errors are moved to the `default` Security Group in AWS and `default-vnet-<vnet-ID>-sg` Network Security Group in Microsoft Azure.
- Unmanaged VMs and VMs in the User Managed list are not affected.

Managing VMs in the NSX Enforced Mode

Follow these steps to start managing successfully onboarded VMs in the NSX Enforced Mode.

Table 28-8. Micro-segmentation workflow for your NSX-managed workload VMs in the NSX Enforced Mode

Task	Instructions
❑ To allow inbound access to workload VMs, create distributed firewall (DFW) rules as required.	See Default Connectivity Strategy for NSX-Managed Workload VMs in the NSX Enforced Mode.
❑ Group your workload VMs using public cloud tags or NSX tags and set up micro-segmentation.	See Set up Micro-segmentation for Workload VMs in the NSX Enforced Mode. See also: Group VMs using NSX and Public Cloud Tags

Default Connectivity Strategy for NSX-Managed Workload VMs in the NSX Enforced Mode

When you deploy the PCG on your Transit VPC/VNet or when you link a Compute VPC/VNet to a Transit, NSX Cloud creates default Security Policies and DFW rules therein for NSX-managed workload VMs.

The two stateless rules are for DHCP access and they do not affect access to your workload VMs.

The two stateful rules are as follows:

DFW Rules created by NSX Cloud under Policy: `cloud-stateful-cloud-<VPC/VNet ID>`	Properties
`cloud-<VPC/VNet ID>-managed`	Allows access to the VMs within the same VPC/VNet.
`cloud-<VPC/VNet ID>-inbound`	Blocks access to NSX-managed VMs from anywhere outside the VPC/VNet.

Note Do not edit any of the default rules.

You can create a copy of the existing inbound rule, adjust the sources and destinations, and set to **Allow**. Place the **Allow** rule above the default **Reject** rule. You can also add new policies and rules. See Add a Distributed Firewall for instructions.

Set up Micro-segmentation for Workload VMs in the NSX Enforced Mode

You can set up micro-segmentation for managed workload VMs.

Note DFW rules depend on the tags assigned to VMs. Since these tags can be modified by anyone with the appropriate public cloud permissions, NSX assumes that such users are trustworthy and the responsibility of ensuring and auditing that VMs are correctly tagged at all times lies with the public cloud network administrator.

Do the following to apply distributed firewall rules to NSX-managed workload VMs:

1 Create groups using VM names or tags or other membership criteria, for example, for **web**, **app**, **DB** tiers. For instructions, see Add a Group.

   You can use any of the following tags for membership criteria. See Group VMs using NSX and Public Cloud Tags for details.

   - system-defined tags
   - tags from your VPC or VNet that are discovered by NSX Cloud
   - or your own custom tags

2 Create an East-West distributed firewall policy and rule and apply to the group you created. See Add a Distributed Firewall . You can also use Context Profiles to create rules specific to App IDs and FQDN/URLs. A predefined list of public cloud FQDN/URLs is available when you create an FQDN/URL context profile. See Layer 7 Context Profile for details.

This micro-segmentation takes effect when the inventory is either manually re-synchronized from CSM, or within about three minutes when the changes are pulled into CSM from your public cloud.

Native Cloud Enforced Mode

In the Native Cloud Enforced Mode, all your workload VMs are automatically NSX-managed. Follow the workflow outlined here to start managing these VMs using NSX.

Note All operating systems are supported for your workload VMs in the Native Cloud Enforced Mode.

Managing VMs in the Native Cloud Enforced Mode

In the Native Cloud Enforced Mode, NSX Cloud utilizes NSX Groups and Distributed Firewall rules to create corresponding Application Security Groups and Network Security Groups in Microsoft Azure and Security Groups in AWS.

All workload VMs in your VPCs/VNets onboarded in the Native Cloud Enforced Mode are NSX-managed.

Follow this workflow:

Table 28-9. Micro-segmentation workflow for your workload VMs in the Native Cloud Enforced Mode

Task	Instructions
❑Create one or more Groups in NSX Manager to include workload VMs from your public cloud.	See Set up Micro-segmentation for Workload VMs in the Native Cloud Enforced Mode See also: Group VMs using NSX and Public Cloud Tags
❑Create one or more Security Policies in NSX Manager that apply to the Group(s) you created for your public cloud workload VMs.	
❑Remove workload VMs from the User Managed list in CSM if you want them managed by NSX Security Policies.	
❑Resync your public cloud account in CSM.	
❑From your VPC/VNet, switch to the details view in CSM for troubleshooting Security policies if there are any errors.	See Current Limitations and Common Errors

Set up Micro-segmentation for Workload VMs in the Native Cloud Enforced Mode

You can configure Security Policy in NSX Manager for workload VMs in the Native Cloud Enforced Mode.

Starting in NSX 3.0, you can create security policies and rules in VPCs/VNets from different accounts or subscriptions.

Note DFW rules depend on the tags assigned to VMs. Since these tags can be modified by anyone with the appropriate public cloud permissions, NSX assumes that such users are trustworthy and the responsibility of ensuring and auditing that VMs are correctly tagged at all times lies with the public cloud network administrator.

Prerequisites

Verify that you have a Transit or Compute VPC/VNet in the Native Cloud Enforced Mode.

Procedure

1 In NSX Manager, edit or create Groups for workload VMs, for example, VM names starting with web, app, db, could be three separate Groups. See Add a Group for instructions. Also see Group VMs using NSX and Public Cloud Tags for information on using public cloud tags to create Groups for your workload VMs.

   Workload VMs that match the criteria are be added to the Group. VMs that do not match any grouping criteria are placed in the `default` Security Group in AWS and the `default-vnet-<vnet-ID>-sg` Network Security Group in Microsoft Azure.

   Note You cannot use the Groups that are auto-created by NSX Cloud.

2 In NSX Manager, create Distributed Firewall (DFW) rules with these Groups in the **Source**, **Destination** or **Applied To** fields. See Add a Distributed Firewall for instructions.

   Note Only Stateful policies are supported for public cloud workload VMs. Stateless policies can be created in NSX Manager but they will not be matched with any Groups that contain your public cloud workload VMs.

   L7 Context Profiles are not supported for DFW rules for workload VMs in the Native Cloud Enforced Mode.

3 In CSM, remove those VMs from the User Managed list that you want to bring under NSX management. See How to use the User Managed List for instructions.

   Note Adding VMs to the User Managed list is a manual step that is strongly recommended in the day-0 workflow as soon as you add your public cloud inventory in CSM. If you have not added any VMs to the User Managed list, you do not need to remove them from it.

4 For Groups and DFW rules that find a match in the public cloud, the following takes place automatically:

   a In AWS, NSX Cloud creates a new Security Group named like `nsx-<NSX GUID>`.

   b In Microsoft Azure, NSX Cloud creates an Application Security Group (ASG) corresponding with the Group created in NSX Manager and a Network Security Group (NSG) corresponding to the DFW rules that are matched with grouped workload VMs.

      NSX Cloud synchronizes NSX Manager and public cloud groups and DFW rules every 30 seconds.

5 Resynchronize your public cloud account in CSM:

   a Log in to CSM and go to your public cloud account.

   b From the public cloud account, click **Actions > Resync Account**. Wait for the resynch to complete.

c Go to the VPC/VNet and click on the red-colored **Errors** indicator. This takes you to the instances view.

d Switch the view to Details if viewing in Grid and click on **Failed** in the Rules Realization column to view errors, if any.

What to do next

See Current Limitations and Common Errors.

Current Limitations and Common Errors

Refer to these known limitations and common errors to troubleshoot managing your public cloud workload VMs in the Native Cloud Enforced Mode.

Note The following limits are set by your public cloud:

- The number of security groups that can be applied to a workload VM.
- The number of rules that can be realized for a workload VM.
- The number of rules that can be realized per security group.
- The scope of the security group assignment, for example, the scope of the Network Security Group (NSG) in Microsoft Azure is limited to that region, whereas the scope of the Security Group (SG) in AWS is limited to that VPC.

Refer to the public cloud documentation for more information on these limits.

Current Limitations

The current release has the following limitations for DFW rules for workload VMs:

- Nested Groups are not supported.
- Groups without VM and/or IP address as member are not supported, for example, Segment or Logical Port based criteria are not supported.
- Both Source and Destination as IP address or CIDR based Group is not supported.
- Both Source and Destination as "ANY" is not supported.
- **Applied_To** Group can be only Source or Destination or Source + Destination Groups. Other options are not supported.

- Only TCP, UDP, and ICMP are supported.

Note Only in AWS:

Deny rules created for workload VMs in your AWS VPCs are not realized on AWS because in AWS, everything is in the denied list by default. This leads to the following results in NSX:

- If there is a Deny rule between VM1 and VM2 then traffic is not allowed between VM1 and VM2 because of the default AWS behavior, not because of the Deny rule. The Deny rule is not realized in AWS.
- Assuming the following two rules are created in NSX Manager for the same VMs with rule 1 having a higher priority than rule 2:

  a VM1 to VM2 DENY SSH

  b VM1 to VM2 Allow SSH

  the Deny rule is ignored because it is not realized in AWS and therefore the Allow SSH rule is realized. This is contrary to expectation but a limitation because of the default AWS behavior.

Common Errors and their Resolution

Error: No NSX policy applied to VM.

If you see this error, none of the DFW rules were applied to the particular VM. Edit the rule or the Group in NSX Manager to include this VM.

Error: Stateless NSX rule is not supported.

If you see this error, it means that you have added DFW rules for public cloud workload VMs in a Stateless Security Policy. This is not supported. Create a new or use an existing Security Policy in the Stateful mode.

NSX Features Supported with NSX Cloud

NSX Cloud creates a network topology for your public cloud VPC or VNet by generating logical networking entities in NSX.

Use this list as a reference for what is auto-generated and how you should use NSX features as they apply to the public cloud.

NSX Manager Configurations

See **Auto-created NSX Logical Entities** in the *NSX Installation Guide* for details on the logical entities created after a PCG is successfully deployed.

Important Do not edit or delete any of these auto-created entities.

Note If you are not able to access some features on Windows workload VMs ensure that the Windows firewall settings are correctly configured.

Table 28-10.

NSX Feature	Details	NSX Cloud Note
Segments or Logical Switches	See Chapter 4 Segments.	A segment is created for every public cloud subnet to which a managed VM is attached. This is a hybrid segment.
Gateways or Logical Routers	See Chapter 2 Tier-0 Gateways and Chapter 3 Tier-1 Gateway.	When PCG is deployed on a Transit VPC or VNet, a tier-0 logical router is auto-created by NSX Cloud. A tier-1 router is created for each Compute VPC/VNet when it's linked to a Transit VPC/VNet
IPFIX	See Configure IPFIX in Manager Mode.	■ IPFIX is supported in NSX Cloud only on UDP port 4739. ■ **Switch and DFW IPFIX**: If the collector is in the same subnet as the Windows VM on which IPFIX profile has been applied, a static ARP entry for the collector on the Windows VM is needed because Windows silently discards UDP packets when no ARP entry is found.
Port Mirroring	See Monitor Port Mirroring Sessions in Manager Mode.	Port Mirroring is supported only in AWS in the current release. ■ For NSX Cloud, configure Port Mirroring from **Tools > Port Mirroring Session**. ■ Only L3SPAN Port Mirroring is supported. ■ The collector must be in the same VPC as the source workload VM.
Distributed Firewall (DFW)	See Distributed Firewall .	■ Layer 4 - Layer 7 with Application IDs. ■ FQDN Filtering.
Gateway Firewall (GFW)	See Gateway Firewall.	Supported on tier-0 gateways.

Group VMs using NSX and Public Cloud Tags

NSX Cloud allows you to use the public cloud tags assigned to your workload VMs.

NSX Manager uses tags to group VMs, as do public clouds. Therefore, to facilitate grouping VMs, NSX Cloud pulls in the public cloud tags applied to your workload VMs provided they meet predefined size and reserved-words criteria, into NSX Manager.

Note DFW rules depend on the tags assigned to VMs. Since these tags can be modified by anyone with the appropriate public cloud permissions, NSX assumes that such users are trustworthy and the responsibility of ensuring and auditing that VMs are correctly tagged at all times lies with the public cloud network administrator.

Tags terminology

A **tag** in NSX Manager refers to what is known as **value** in a public cloud context. The **key** of a public cloud tag, is referred to as **scope** in NSX Manager.

Components of tags in NSX Manager	Equivalent components of tags in the public cloud
Scope	Key
Tag	Value

Tag Types and Limitations

NSX Cloud allows three types of tags for NSX-managed public cloud VMs.

- **System Tags**: These tags are system-defined and you cannot add, edit, or delete them. NSX Cloud uses the following system tags:
  - azure:subscription_id
  - azure:region
  - azure:vm_rg
  - azure:vnet_name
  - azure:vnet_rg
  - azure:transit_vnet_name
  - azure:transit_vnet_rg
  - aws:account
  - aws:availabilityzone
  - aws:region
  - aws:vpc
  - aws:subnet
  - aws:transit_vpc

- **Discovered Tags**: Tags that you have added to your VMs in the public cloud are automatically discovered by NSX Cloud and displayed for your workload VMs in NSX Manager inventory. These tags are not editable from within NSX Manager. There is no limit to the number of discovered tags. These tags are prefixed with `dis:azure:` to denote they are discovered from Microsoft Azure and `dis:aws` from AWS.

  When you make any changes to the tags in the public cloud, the changes are reflected in NSX Manager within three minutes.

  By default this feature is enabled. You can enable or disable the discovery of Microsoft Azure or AWS tags at the time of adding the Microsoft Azure subscription or AWS account.

- **User Tags**: You can create up to 25 user tags. You have add, edit, delete privileges for user tags. For information on managing user tags, see Manage Tags for a VM in Manager Mode.

Table 28-11. Summary of Tag Types and Limitations

Tag type	Tag scope or predetermined prefix	Limitations	Enterprise Administrator Privileges	Auditor Privileges
System-defined	Complete system tags: ■ azure:subscription_id ■ azure:region ■ azure:vm_rg ■ azure:vnet_name ■ azure:vnet_rg ■ aws:vpc ■ aws:availability zone	Scope (key): 20 characters Tag (value): 65 characters Maximum possible: 5	Read only	Read only
Discovered	Prefix for Microsoft Azure tags that are imported from your VNet: **dis:azure:** Prefix for AWS tags that are imported from your VPC: **dis:aws:**	Scope (key): 20 characters Tag (value): 65 characters Maximum allowed: unlimited **Note** The limits on characters excludes the prefix **dis:<public cloud name>**. Tags that exceed these limits are not reflected in NSX Manager. Tags with the prefix **nsx** are ignored.	Read only	Read only
User	User tags can have any scope (key) and value within the allowed number of characters, except: ■ the scope (key) prefix **dis:azure:** or **dis:aws:** ■ the same scope (key) as system tags	Scope (key): 30 characters Tag (value): 65 characters Maximum allowed: 25	Add/Edit/Delete	Read only

Examples of Discovered Tags

Note Tags are in the format **`key=value`** for the public cloud and **`scope=tag`** in NSX Manager.

Table 28-12.

Public Cloud tag for the workload VM	Discovered by NSX Cloud?	Equivalent NSX Manager tag for the workload VM
Name=Developer	Yes	dis:azure:Name=Developer
ValidDisTagKeyLength=ValidDisTagValue	Yes	dis:azure:ValidDisTagKeyLength=ValidDisTagValue
Abcdefghijklmnopqrstuvwxyz=value2	No (key exceeds 20 chars)	none
tag3=AbcdefghijklmnopqrstuvwxyzAb23690hgjgjuytreswqacvbcdefghijklmnopqrstuvwxyz	No (value exceeds 65 characters)	none
nsx.name=Tester	No (key has the prefix **nsx**)	none

How to use Tags in NSX Manager

- See Manage Tags for a VM in Manager Mode.
- See Search for Objects.
- See Add a Group.
- See Set up Micro-segmentation for Workload VMs in the NSX Enforced Mode.

Use Native-Cloud Services

The following native-cloud services are supported for use with your public cloud workload VMs from within NSX Manager.

When you deploy PCG, a Group is created in NSX Manager for each supported native-cloud service.

The following Groups are created for the currently supported public cloud services:

- `aws-dynamo-db-service-endpoint`
- `aws-elb-service-endpoint`
- `aws-rds-service-endpoint`
- `aws-s3-service-endpoint`
- `azure-cosmos-db-service-endpoint`
- `azure-load-balancer-service-endpoint`
- `azure-sql-service-endpoint`
- `azure-storage-service-endpoint`

To use these native-cloud services, create DFW policies that contain the native-cloud service Group in the Source or Destination fields of the rule as required.

DFW rules are enforced on VMs not on the native-cloud services.

Note In the NSX Enforced Mode, that is, managing your workloads with NSX Tools, currently there is no support for Microsoft Azure's native-cloud services.

Current Limitations

ENDPOINT			DFW Rule with service as DESTINATION		DFW Rule with service as SOURCE	
Public Cloud	**Service**	**Scope**	**Enforced on VM?**	**Enforced on Service?**	**Enforced on Service?**	**Enforced on VM?**
Microsoft Azure	BLOB Storage	Global	Yes	No	No	Yes
	Cosmos DB					
	SQL					
	Load Balancer					
AWS	S3	VPC Local	Yes	No	No	Yes
	Dynamo DB					
	RDS					
	ELB					

Service Insertion for your Workload VMs in the NSX Enforced Mode

NSX Cloud supports the use of third-party services in your public cloud for NSX-managed workload VMs in the NSX Enforced Mode.

NSX Cloud supports Service Insertion for the following:

- North-south traffic from workload VMs via a service appliance hosted in a Transit VPC/VNet..
- VPN traffic from the PCG to an on-prem edge or gateway. This traffic can be routed via a service appliance in a Transit VPC/VNet as well.

Here is an overview of the configurations to allow service insertion for your NSX-managed workload VMs.

Table 28-13. Overview of configurations required for service insertion for NSX-managed workload VMs in the NSX Enforced Mode.

Frequency	Task	Instructions
Follow these instructions for the initial setup if you want to set up service insertion for north-south traffic.	Set up the service appliance in your public cloud preferably in a Transit VPC or VNet (where you have deployed the PCG.	See instructions specific to the third-party service appliance and the public cloud.
	Register the third-party service in NSX.	See Create the Service Definition and a Corresponding Virtual Endpoint
	Create a virtual instance endpoint of the service using a /32 Virtual Service IP address (VSIP) to be used only for service insertion by the service appliance. The VSIP should not conflict with the CIDR range of VPCs or VNets. This VSIP is advertised over BGP to the PCG.	See Create the Service Definition and a Corresponding Virtual Endpoint
	Create an IPSec VPN tunnel between the service appliance and the PCG.	See Set up an IPSec VPN Session
	Configure BGP between the PCG and the service appliance and advertise the VSIP from the service appliance and the default route (0.0.0.0/0) from the PCG.	See Configure BGP and Route Redistribution
Follow these instructions for the initial setup for VPN traffic from the public cloud to on-prem.	Create a VPN tunnel between the PCG and the on-prem edge or gateway.	See Set up VPN in the NSX Enforced Mode.
Follow these instructions for both types of service insertion as part of the initial setup.	Create a lowest priority default catch-all rule with the action set to **Do Not Redirect**. This ensures that no packets are redirected on the VTI interface of the PCG and the Service Appliance.	See Set up Redirection Rules.
Follow these instructions as and when necessary for each type of service insertion use case.	After the one-time configurations are complete, set up redirection rules to reroute selective traffic from NSX-managed workload VMs to the VSIP. These rules are applied to the uplink port of the PCG for north-south service insertion and to the VTI interface of the PCG for traffic to on-prem.	See Set up Redirection Rules.

Procedure

1 Create the Service Definition and a Corresponding Virtual Endpoint

You must use NSX Manager APIs to create a service definition and virtual endpoint for the service appliance in your public cloud.

2 Set up an IPSec VPN Session

Set up an IPSec VPN session between the PCG and your service appliance.

3 Configure BGP and Route Redistribution

Configure BGP between the PCG and the service appliance over the IPSec VPN tunnel.

4 Set up Redirection Rules

You must set up a default redirection rule as part of the initial setup for service insertion.

Create the Service Definition and a Corresponding Virtual Endpoint

You must use NSX Manager APIs to create a service definition and virtual endpoint for the service appliance in your public cloud.

Prerequisites

Pick out a /32 reserved IP address to serve as the Virtual Endpoint for the service appliance in your public cloud, for example, `100.100.100.100/32`. This is referred to as the Virtual Service IP (VSIP).

Note If you deployed your service appliance in a High Availability pair, do not create another service definition but use the same VSIP when advertising it to the PCG during BGP configuration.

Procedure

1 To create a Service Definition for the service appliance, run the following API call using NSX Manager credentials for authorization:

```
POST https://{{NSX Manager-IP}}/policy/api/v1/enforcement-points/default/service-definitions
```

Example request:

```
{
     "resource_type":"ServiceDefinition",
     "description":"NS-Service",
     "display_name":"Service_Appliance1",
     "attachment_point":[
        "TIER0_LR"
     ],
     "transports":[
        "L3_ROUTED"
     ],
     "functionalities":[
        "NG_FW", "BYOD"
     ],
     "on_failure_policy":"ALLOW",
     "implementations":[
```

```
        "NORTH_SOUTH"
     ],
     "vendor_id" : "Vendor1"
}
```

Example response:

```
{
    "resource_type": "ServiceDefinition",
    "description": "NS-Service",
    "id": "33890153-6eea-4c9d-8e34-7b6532b9d65c",
    "display_name": "Service_Appliance1",
    "attachment_point": [
        "TIER0_LR"
    ],
    "transports": [
        "L3_ROUTED"
    ],
    "functionalities": [
        "NG_FW", "BYOD"
    ],
    "vendor_id": "Vendor1",
    "on_failure_policy": "ALLOW",
    "implementations": [
        "NORTH_SOUTH"
    ],
    "_create_time": 1540424262137,
    "_last_modified_user": "nsx_policy",
    "_system_owned": false,
    "_protection": "REQUIRE_OVERRIDE",
    "_last_modified_time": 1540424262137,
    "_create_user": "nsx_policy",
    "_revision": 0
}
```

2 To create a Virtual Endpoint for the service appliance, run the following API call using NSX Manager credentials for authorization:

```
PATCH https://{{NSX Manager-IP}}policy/api/v1/infra/tier-0s/<tier-0 router ID>/locale-
services/cloud/endpoints/virtual-endpoints/Service_Appliance1_Endpoint
```

Example request:

```
{
  "resource_type": "VirtualEndpoint",
  "display_name": "Service_Appliance1_Endpoint",
  "target_ips": [
    {
      "ip_addresses": [
        "100.100.100.100"
      ],
      "prefix_length": 32
    }
  ],
```

```
  "service_names": [
    "Service_Appliance1"
  ]
}
```

Example response:

```
200 OK
```

Note The `display_name` in step 1 must match the `service_names` in step 2.

What to do next

Set up an IPSec VPN Session

Set up an IPSec VPN Session

Set up an IPSec VPN session between the PCG and your service appliance.

Prerequisites

- One or an HA pair of PCGs must be deployed in a Transit VPC/VNet.
- The service appliance must be set up in your public cloud, preferably in the Transit VPC/VNet.

Procedure

1. Navigate to **Networking > VPN**
2. Add a **VPN service** of type IPSec and note the following configuration options specific to NSX Cloud. See Add an IPSec VPN Service for other details.

Option	Description
Name	The name of this VPN service is used to set up the local endpoint and the IPSec VPN sessions. Make a note of it.
Service Type	Confirm that this value is set to IPSec.
Tier-0 Gateway	Select the tier-0 gateway auto-created for your Transit VPC/VNet. Its name contains your VPC/VNet ID, for example, `cloud-t0-vpc-6bcd2c13`.

3 Add a **Local Endpoint** for your PCG. The IP address of the local endpoint is the value of the tag `nsx:local_endpoint_ip` for the PCG deployed in your Transit VPC/VNet. Log in to your Transit VPC/VNet for this value. Note the following configurations specific to NSX Cloud and see Add Local Endpoints for other details.

Option	Description
Name	The local endpoint name is used to set up the IPSec VPN sessions. Make a note of it.
VPN Service	Select the VPN Service you added in step 2.
IP Address	Find this value by logging in to the AWS console or the Microsoft Azure portal. It is the value of the tag `nsx:local_endpoint_ip` applied to the uplink interface of the PCG.

4 Create a **Route-Based IPSec session** between the PCG and the service appliance in your public cloud (preferably hosted in the Transit VPC/VNet).

Option	Description
Type	Confirm that this value is set to **Route Based**.
VPN Service	Select the VPN Service you added in step 2.
Local Endpoint	Select the local endpoint you created in step 3.
Remote IP	Enter the private IP address of the service appliance. **Note** If your service appliance is accessible using a public IP address, assign a public IP address to the local endpoint IP (also known as secondary IP) to the PCG's uplink interface.
Tunnel Interface	This subnet must match with the service appliance subnet for the VPN tunnel. Enter the subnet value you set up in the service appliance for the VPN tunnel or note the value you enter here and make sure the same subnet is used when setting up the VPN tunnel in the service appliance. **Note** You configure BGP on this tunnel interface. See Configure BGP and Route Redistribution .
Remote ID	Enter the private IP address of your service appliance in the public cloud.
IKE Profile	The IPSec VPN session must be associated with an IKE profile. If you created a profile, select it from the drop-down menu. You can also use the default profile.

What to do next

Configure BGP and Route Redistribution

Configure BGP and Route Redistribution

Configure BGP between the PCG and the service appliance over the IPSec VPN tunnel.

You set up BGP neighbors on the IPSec VPN tunnel interface that you established between PCG and the service appliance. See Configure BGP for more details.

You need to configure BGP similarly on your service appliance. See documentation for your specific service in the public cloud for details.

Next, set up route redistribution as follows:

- The PCG advertises its default route (0.0.0.0/0) to the service appliance.
- The service appliance advertises the VSIP to the PCG. This is the same IP address which is used when registering the service. See Create the Service Definition and a Corresponding Virtual Endpoint.

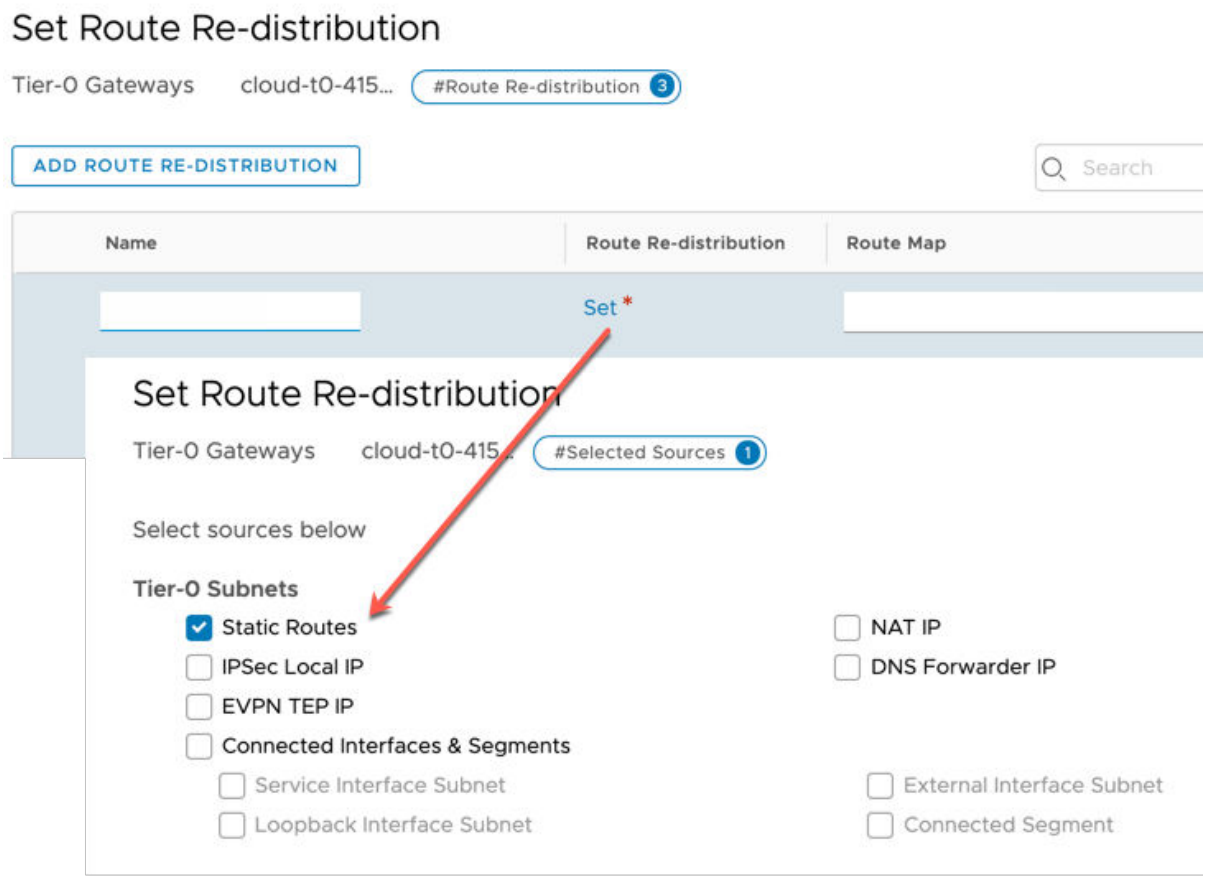
Note If your service appliance is deployed in a High Availability pair, advertise the same VSIP from both service appliances.

Procedure

1. Navigate to **Networking > Tier-0 Gateways** .
2. Select the auto-created tier-0 gateway for your Transit VPC/VNet named like `cloud-t0-vpc-6bcd2c13` and click **Edit**.
3. Click the number or icon next to **BGP Neighbors** under the **BGP** section.
4. Note these configurations:

Option	Description
IP Address	Use the IP address configured on the service appliance tunnel interface for the VPN between the PCG and the service appliance.
Remote AS Number	This number must match the AS number of the service appliance in your public cloud.
Route Filter	Set an Out Filter to advertise the default route (0.0.0.0/0) from the PCG to service appliance.

5 From the **Route Redistribution** section, enable static routes on tier-0 gateway.



What to do next

Set up Redirection Rules

Set up Redirection Rules

You must set up a default redirection rule as part of the initial setup for service insertion.

After the initial setup is completed, you can create and edit redirection rules as necessary for rerouting different types of traffic for your NSX-managed workload VMs through the service appliance.

These are the two types of redirection rules:

1 As part of the initial service insertion setup, you must create a catch-all rule to prevent redirection for traffic for the VTI interface of the VPN tunnel between the PCG and the service appliance. This rule must have the lowest possible priority and must be created for both use cases of service insertion.

2 The second rule sets up specific redirection for traffic for the service appliance. You can adjust this rule and add others as necessary.

Procedure

1 To add the default catch-all rule to complete the one-time setup, follow these steps:

a Navigate to **Security > North South Firewall > Network Introspection (N-S)**

b Click **Add Policy**.

Option	Description
Name	Provide a descriptive name, for example, **`Default_No-Redirect-Policy`**.
Redirect To:	Select the name of the Virtual Endpoint you created for this service appliance when registering the service.
Apply To:	Select the PCG's tier-0 gateway.

c Select the new policy and click **Add Rule**. Note the following values specific to service insertion:

Option	Description
Sources	Any
Destinations	Any
Applied To	Select the VTI interface between the PCG and the service appliance.
Action	Select **Do Not Redirect**.

Important This rule must have the lowest possible priority.

2 For the second rule, follow these steps:

a Navigate to **Security > North South Firewall > Network Introspection (N-S)**

b Click **Add Policy**.

Option	Description
Name:	Provide a descriptive name for the policy, for example, **`On-Prem Service Insertion for AWS VMs`** or **`North-south Service Insertion for Azure VMs`**.
Redirect To:	Select the name of the Virtual Endpoint you created for this service appliance when registering the service.
Apply To:	Select the PCG's tier-0 gateway.

c Select the new policy and click **Add Rule**. Note the following values specific to service insertion:

Option	Description
Sources	Select a group of subnets whose traffic must be redirected, for example, a group of your NSX-managed workload VMs.
Destinations	Select a list of destination IP addresses or services, such as **`YouTube`**, that you want to route through the service appliance.
Applied To	■ If you are using north-south service insertion with the service appliance in the public cloud: select the uplink port of the active and standby PCG. ■ If you are using VPN traffic to on-prem: select the VTI interface of the active and standby PCG to the on-prem service appliance.
Action	Select **Redirect**.

Enable NAT on NSX-managed VMs

NSX Cloud supports enabling NAT on NSX-managed VMs.

You can enable North-South traffic on VMs in NSX-managed VMs using public cloud tags.

On the NSX-managed VM for which you want to enable NAT, apply the following tag:

Table 28-14.

Key	Value
`nsx.publicip`	**`public IP address from your public cloud`**, for example, 50.1.2.3

Note The public IP address you provide here must be free to use and must not be assigned to any VM, even the workload VM you want to enable NAT for. If you assign a public IP address that was previously associated with any other instance or private IP address, NAT does not work. In that case, unassign the public IP address.

After this tag is applied, the workload VM can access internet traffic.

Enable Syslog Forwarding

NSX Cloud supports syslog forwarding.

You can enable syslog forwarding for Distributed Firewall (DFW) packets on managed VMs.

To learn more about how to configure NSX appliances and hypervisors to send log messages to a remote logging server, see Configure Remote Logging.

If logs are not received by the remote log server, see Troubleshooting Syslog Issues.

Procedure

1. Log in to PCG using the jump host.
2. Type **`nsxcli`** to open the NSX CLI.
3. Type the following commands to enable DFW log forwarding:

   ```
   nsx-public-cloud-gateway> set gw-controller vm-log-forwarding enabled
   nsx-public-cloud-gateway> set logging-server <server-IP-address> proto udp level info
   messageid FIREWALL-PKTLOG
   ```

   After this is set, NSX agent DFW packet logs are available under `/var/log/syslog` on PCG.
4. To enable log forwarding per VM, enter the following command:

   ```
   nsx-public-cloud-gateway> set gw-controller vm-log-forwarding enabled <vm-id>
   ```

Automate VPN for Public Cloud Endpoints using APIs

You can use CSM APIs to automate VPN setup between VPCs or VNets.

You cannot use CSM APIs to set up VPN using on-prem endpoints.

Prerequisites:

- Both endpoints for the VPN tunnel must be in the public cloud with PCGs deployed in them and in a running or `up` state.

The following entities are configured using CSM APIs. You can also use these APIs to unconfigure the VPN setup.

For NSX entities that support tags, the API reuses previously created entities by applying a tag to them with the new VPN session ID, for example, `CsmVpnSessionId:<csm-vpn-session-id-new>`.

- For each source and destination endpoint:
  - IPSec VPN service is configured, named `cloud-vpn-service-<vpc/vnet-id>`.
  - BGP routing is configured, named `cloud-routing-config-<vpc/vnet-id>`. If BGP was not already enabled, the API enables it and assigns an AS number in the format: `55555.<1-64999>`.

- For each PCG in source and destination endpoints:
  - BGP route re-distribution is enabled for tier-0 static routes and tier-1 connected segments.
  - Public IP is assigned to PCG's uplink interface and associated to VPN-secondary private IP on PCG's uplink interface.
  - IPSec VPN local endpoint is created, named `cloud-vpn-local-endpoint-<gateway-id>-<preferred/non-preferred>`.
- For each PCG combination between source and destination endpoints:
  - Route based IPSec VPN session is created, named `<csm-vpn-session-id>-<Preferred/non-preferred>To<Preferred/non-preferred>-<hash-from-source-and-destination-PCG-ids>`
  - BGP neighbor is added on tier-0 gateway for each IPSec VPN session configured.

Configuring/Updating VPN Sessions

1 To configure a new VPN session, do the following:

a
```
GET /api/v1/csm/vpn/endpoints
```

b
```
POST /api/v1/csm/vpn/session

Example Request:
POST https://<nsx-csm>/api/v1/csm/vpn/sessions
{
  "display_name": "aws azure session 01",
  "source_endpoint": {
    "id": "vpc-12345678",
    "display_name": "vpc test",
    "endpoint_type": "AWS"
  },
  "destination_endpoint": {
    "id": "d02af61a-e212-486e-b6c8-10462ccfbad6",
    "display_name": "vnet-01",
    "endpoint_type": "AZURE"
  }
}
```

2 To update the display name of an existing VPN session:

```
PUT /api/v1/csm/vpn/sessions/<session-id>

Example Request:
PUT https://<nsx-csm>/api/v1/csm/vpn/sessions/9174ffd1-41b1-42d6-a28d-05c61a0698e2
{
  "display_name": "New VPN session",
  "source_endpoint": {
    "id": "vpc-12345678",
    "display_name": "vpc test",
```

```
    "endpoint_type": "AWS"
  },
  "destination_endpoint": {
    "id": "d02af61a-e212-486e-b6c8-10462ccfbad6",
    "display_name": "vnet-01",
    "endpoint_type": "AZURE"
  }
}
```

Getting the status of existing VPN sessions

- To get status of all sessions:

```
GET /api/v1/csm/vpn/sessions/status
```

- To get the status of a specific session by providing the session-id:

```
GET /api/v1/csm/vpn/sessions/<session-id>/status
```

Deleting Sessions

Delete sessions by providing session-id:

```
DELETE /api/v1/csm/vpn/sessions/<session-id>
```

Troubleshooting

If the creation of entities fails:

- Get the status for the specific session-id:

```
GET /api/v1/csm/vpn/sessions/<session-id>/status
```

- You can see the point of failure in the response. Make the necessary changes to resolve the failure.
- Recreate the remaining entities for the same session id using the API call:

```
POST /api/v1/csm/vpn/sessions/<session-id>?action=recreate
```

See the latest version of the *NSX REST API Guide* at https://code.vmware.com/ for API details.

Set up VPN in the Native Cloud Enforced Mode

You can create a VPN tunnel between the PCG and a remote endpoint by following this workflow. These instructions are specific to workload VMs manged in the Native Cloud Enforced Mode.

You can use CSM APIs to configure VPN in NSX if both the endpoints are in the public cloud and managed by PCGs. See Automate VPN for Public Cloud Endpoints using APIs.

Prerequisites

- In AWS: Verify that you have deployed a VPC in the Native Cloud Enforced Mode. This must be a Transit or Self-managed VPC. VPN is not supported for Compute VPCs in AWS.
- In Microsoft Azure: Verify that you have deployed a VNet in the Native Cloud Enforced Mode. You can use both Transit and Compute VNets.
- Verify that the remote endpoint is peered with the PCG and has route-based IPSec VPN and BGP capabilities.

Procedure

1. In your public cloud, find the NSX-assigned local endpoint for the PCG and assign a public IP address to if necessary:
   a. Go to your PCG instance in the public cloud and navigate to Tags.
   b. Note the IP address in the value field of the tag `nsx.local_endpoint_ip`.
   c. (Optional) If your VPN tunnel requires a public IP, for example, if you want to set up a VPN to another public cloud or to the on-prem NSX deployment:
      1. Navigate to the uplink interface of the PCG instance.
      2. Attach a public IP address to the `nsx.local_endpoint_ip` IP address that you noted in step 1.b.
   d. (Optional) If you have an HA pair of PCG instances, repeat steps 1.a and 1.b and attach a public IP address if necessary, as described in step 1.c.

2 In NSX Manager, enable IPSec VPN for the PCG that appears as a tier-0 gateway named like `cloud-t0-vpc/vnet-<vpc/vnet-id>` and create route-base IPSec sessions between this tier-0 gateway's endpoint and the remote IP address of the desired VPN peer. See Add an IPSec VPN Service for other details.

a Navigate to **Networking > VPN > VPN Services > Add Service > IPSec**. Provide the following details:

Option	Description
Name	Enter a descriptive name for the VPN service, for example **<VPC-ID>-AWS_VPN** or **<VNet-ID>-AZURE_VPN**.
Tier0/Tier1 Gateway	Select the tier-0 gateway for the PCG in your public cloud.

b Navigate to **Networking > VPN > Local Endpoints > Add Local Endpoint**. Provide the following information and see Add Local Endpoints for other details:

Note If you have an HA pair of PCG instances, create a local endpoint for each instance using the corresponding local endpoint IP address attached to it in the public cloud.

Option	Description
Name	Enter a descriptive name for the local endpoint, for example **<VPC-ID>-PCG-preferred-LE** or **<VNET-ID>-PCG-preferred-LE**
VPN Service	Select the VPN service for the PCG's tier-0 gateway that you created in step 2.a.
IP Address	Enter the value of the PCG's local endpoint IP address that you noted in step 1.b.

c Navigate to **Networking > VPN > IPSec Sessions > Add IPSec Session > Route Based**. Provide the following information and see Add a Route-Based IPSec Session for other details:

Note If you are creating a VPN tunnel between PCGs deployed in a VPC and PCGs deployed in a VNet, you must create a tunnel for each PCG's local endpoint in the VPC and the remote IP address of the PCG in the VNet, and conversely from the PCGs in the VNet to the remote IP address of PCGs in the VPC. You must create a separate tunnel for the active and standby PCGs. This results in a full mesh of IPSec Sessions between the two public clouds.

Option	Description
Name	Enter a descriptive name for the IPsec session, for example, **<VPC--ID>-PCG1-to-remote_edge**
VPN Service	Select the VPN service you created in step 2.a.
Local Endpoint	Select the local endpoint you created in step 2.b.
Remote IP	Enter the public IP address of the remote peer with which you are creating the VPN tunnel.

Option	Description
	Note Remote IP can be a private IP address if you are able to reach the private IP address, for example, using DirectConnect or ExpressRoute.
Tunnel Interface	Enter the tunnel interface in a CIDR format. The same subnet must be used for the remote peer to establish the IPSec session.

3 Expand **BGP** and set up BGP neighbors on the IPSec VPN tunnel interface that you established in step 2. See Configure BGP for more details.

a Navigate to **Networking > Tier-0 Gateways**.

b Select the auto-created tier-0 gateway for which you created the IPSec session and click **Edit**.

c Click the number or icon next to **BGP Neighbors** under the **BGP** section and provide the following details:

Option	Description
IP Address	Use the IP address of the remote VTI configured on the tunnel interface in the IPSec session for the VPN peer.
Remote AS Number	This number must match the AS number of the remote peer.

4 Advertise the prefixes you want to use for the VPN using the Redistribution Profile. Do the following:

Important This step is only for NSX 3.0.0. Skip it if you are using NSX 3.0.1.

a Expand **Routing** and add a static route for the CIDR of the VPC/VNet onboarded with the Native Cloud Enforced Mode to point to the uplink IP address of the tier-0 gateway, that is, PCG.

See Configure a Static Route for instructions. If you have a PCG pair for HA, set up next hops to each PCG's uplink IP address.

b In the expanded **Routing** category, add a prefix list for the VPC/VNet CIDR onboarded in the Native Cloud Enforced Mode and add it as an Out Filter in BGP neighbor configuration.

See Create an IP Prefix List for instructions.

c Expand **Route Re-Distribution** and set up a route redistribution profile, enabling static route and select the route filter you created for VPC/VNet CIDRs in the previous substep.

5 In your public cloud, do the following:

a Go to the routing table of the subnet where you have your workload VMs.

Note Do not use the routing table of the PCG's uplink or management subnets.

b Add the tag **`nsx.managed = true`** to the routing table.

Results

Verify that routes are created in the managed routing table for all IP prefixes advertised by the remote endpoint with next hop set to the PCG's uplink IP address.

Set up VPN in the NSX Enforced Mode

You can set up VPN using PCGs that appear as auto-created tier-0 gateways in the on-prem NSX deployment. These instructions are specific to workload VMs manged in the NSX Enforced Mode.

Use PCGs in the same way as you use tier-0 gateways in NSX Manager to set up VPN by following the additional steps outlined here. You can create VPN tunnels between PCGs deployed in the same public cloud, or different public clouds, or with an on-prem gateway or router. See Chapter 7 Virtual Private Network (VPN) for details on VPN support in NSX.

You can use CSM APIs to configure VPN in NSX if both the endpoints are in the public cloud and managed by PCGs. See Automate VPN for Public Cloud Endpoints using APIs.

Prerequisites

- Verify that you have one or an HA pair of PCGs deployed in a VPC/VNet.
- Verify that the remote peer supports route-based VPN and BGP.

Procedure

1. In your public cloud, find the NSX-assigned local endpoint for the PCG and assign a public IP address to if necessary:
   a. Go to your PCG instance in the public cloud and navigate to Tags.
   b. Note the IP address in the value field of the tag `nsx.local_endpoint_ip`.
   c. (Optional) If your VPN tunnel requires a public IP, for example, if you want to set up a VPN to another public cloud or to the on-prem NSX deployment:
      1. Navigate to the uplink interface of the PCG instance.
      2. Attach a public IP address to the `nsx.local_endpoint_ip` IP address that you noted in step 1.b.
   d. (Optional) If you have an HA pair of PCG instances, repeat steps 1.a and 1.b and attach a public IP address if necessary, as described in step 1.c.

2 In NSX Manager, enable IPSec VPN for the PCG that appears as a tier-0 gateway named like `cloud-t0-vpc/vnet-<vpc/vnet-id>` and create route-base IPSec sessions between this tier-0 gateway's endpoint and the remote IP address of the desired VPN peer. See Add an IPSec VPN Service for other details.

a Navigate to **Networking > VPN > VPN Services > Add Service > IPSec**. Provide the following details:

Option	Description
Name	Enter a descriptive name for the VPN service, for example **<VPC-ID>-AWS_VPN** or **<VNet-ID>-AZURE_VPN**.
Tier0/Tier1 Gateway	Select the tier-0 gateway for the PCG in your public cloud.

b Navigate to **Networking > VPN > Local Endpoints > Add Local Endpoint**. Provide the following information and see Add Local Endpoints for other details:

Note If you have an HA pair of PCG instances, create a local endpoint for each instance using the corresponding local endpoint IP address attached to it in the public cloud.

Option	Description
Name	Enter a descriptive name for the local endpoint, for example **<VPC-ID>-PCG-preferred-LE** or **<VNET-ID>-PCG-preferred-LE**
VPN Service	Select the VPN service for the PCG's tier-0 gateway that you created in step 2.a.
IP Address	Enter the value of the PCG's local endpoint IP address that you noted in step 1.b.

c Navigate to **Networking > VPN > IPSec Sessions > Add IPSec Session > Route Based**. Provide the following information and see Add a Route-Based IPSec Session for other details:

Note If you are creating a VPN tunnel between PCGs deployed in a VPC and PCGs deployed in a VNet, you must create a tunnel for each PCG's local endpoint in the VPC and the remote IP address of the PCG in the VNet, and conversely from the PCGs in the VNet to the remote IP address of PCGs in the VPC. You must create a separate tunnel for the active and standby PCGs. This results in a full mesh of IPSec Sessions between the two public clouds.

Option	Description
Name	Enter a descriptive name for the IPsec session, for example, **<VPC--ID>-PCG1-to-remote_edge**
VPN Service	Select the VPN service you created in step 2.a.
Local Endpoint	Select the local endpoint you created in step 2.b.
Remote IP	Enter the public IP address of the remote peer with which you are creating the VPN tunnel.

Option	Description
	Note Remote IP can be a private IP address if you are able to reach the private IP address, for example, using DirectConnect or ExpressRoute.
Tunnel Interface	Enter the tunnel interface in a CIDR format. The same subnet must be used for the remote peer to establish the IPSec session.

3 Expand **BGP** and set up BGP neighbors on the IPSec VPN tunnel interface that you established in step 2. See Configure BGP for more details.

a Navigate to **Networking > Tier-0 Gateways**.

b Select the auto-created tier-0 gateway for which you created the IPSec session and click **Edit**.

c Click the number or icon next to **BGP Neighbors** under the **BGP** section and provide the following details:

Option	Description
IP Address	Use the IP address of the remote VTI configured on the tunnel interface in the IPSec session for the VPN peer.
Remote AS Number	This number must match the AS number of the remote peer.

4 Expand **Route Re-Distribution** and advertise the prefixes you want to use for the VPN using the Redistribution Profile. In NSX Enforced Mode, connect tier-1 enabled routes in the redistribution profile.

Deploying NSX Management Components On Microsoft Azure

You can directly deploy the NSX Management appliances natively on the Microsoft Azure.

For more details on architecture and components, refer to the *NSX Installation Guide*.

Task	Instructions
❑ Deployment: You can deploy the NSX Management components using the following two ways: ■ Deploy with Terraform scripts: NSX Cloud provides Terraform scripts that deploys management components –NSX Manager and CSM in your Microsoft Azure subscription. ■ Deploy without Terraform scripts: If you cannot use Terraform scripts, you can also manually deploy all the appliances natively in your Microsoft Azure subscription.	■ Deploy with Terraform scripts ■ Deploy without Terraform scripts
❑ (Optional) Peer VNet to Management VNet (for VGW): As explained in the architecture diagram https://docs.vmware.com/en/VMware-NSX-T-Data-Center/3.2/installation/GUID-2213C06D-C616-4AE7-9AA0-8C8074E779D2.html, you have to deploy the NSX components in the **Management VNet** which is separate from the VNets that you intend to manage. If you do not want use the public IPs and want to ensure that all the communication between the NSX Management components and the PCG+VMs to be private, then you must peer the Management VNet with the VNet that you need to manage. After peering, you can deploy the PCG on the VGW mode and all communications between the two VNets becomes private.	https://docs.vmware.com/en/VMware-NSX-T-Data-Center/3.2/installation/GUID-287768FB-8F52-463D-898A-288B855CCB25.html
❑ Deploy Gateway in the VNet where you want to manage the VMs.	https://docs.vmware.com/en/VMware-NSX-T-Data-Center/3.2/installation/GUID-287768FB-8F52-463D-898A-288B855CCB25.html
❑ Manage your VMs: ■ For NSX Enforced Mode. ■ For Native Cloud Enforced Mode.	■ NSX Enforced Mode ■ Native Cloud Enforced Mode
❑ Manage your backup and restore: ■ Recovery option by redeploying the NSX Manager image. ■ Recovery option using the Microsoft Azure Recovery Vault Service (if you have enabled the vault feature in the Terraform script).	Managing Backup and Restore of NSX Manager and CSM in Microsoft Azure

Redeploying Manager Nodes on Cloud Native Azure

If your NSX appliance gets corrupted or outdated then you must deploy the appliance again. Due to Azure restrictions, you must use the Azure CLI to boot the new VM.

To redeploy the NSX Manager nodes on Azure:

1 Delete the non-functional NSX Manager and its related NIC or disk entities from the Azure resource group.

2 Get the cluster status using the `get cluster status` command through `nsxcli` on any of the remaining NSX Manager nodes.

3 Detach the non-functional NSX Manager node from the NSX Manager cluster. For example,

```
detach node 8992e79f-219f-2c42-be57-c4d576792b78
```

```
Node has been detached. Detached node must be deleted permanently.
```

4 Create a custom data as per step 7 of the https://docs.vmware.com/en/VMware-NSX-T-Data-Center/3.2/installation/GUID-71DDCE82-0F4F-4E75-A117-FB398A1FDFCB.html topic.

5 Create a new manager node using the following command. The following command will add data disk of 100 GB). Change the data disk and release number as per your requirement. Store the public key in the location where you want to run the following command.

```
az vm create --name <MP instance name> --resource-group <RG for MP
deployment> --admin-username nsxadmin --public-ip-address-allocation static --size
Standard_D4_v4 --subnet <subnet_path> --nsg <mgr_nsg_path>  --image vmware-inc:nsx-policy-
manager:byol_release-3-1:3.5.0 --storage-sku Standard_LRS --data-disk-sizes-gb 100 --
authentication-type ssh --ssh-key-values <publickey_path> --custom-data <userdata_txt_path>
```

6 Wait for around 15 minutes for the services and cluster to be up on the new single node cluster.

7 Join this new node to the existing cluster by running the following command on the new NSX Manager node through `nsxcli`.

```
join 192.168.1.11 cluster-id 95e888bf-d8fb-4974-8da7-13029d7be8f0
username nsxadmin password <password> thumbprint
32135bdbd14fe3cba11e1d91b106c2f1e28e0d464c23bbe3caf88fdf44b0eca2
```

```
Data on this node will be lost. Are you sure? (yes/no): yes Join operation
successful. Services are being restarted. Cluster may take some time to stabilize.
```

Wait for around 15 minutes for the three-node cluster to be up and running.

Managing Backup and Restore of NSX Manager and CSM in Microsoft Azure

You can use Microsoft Azure services to restore NSX Manager and CSM appliances that are deployed natively in Microsoft Azure.

Note This information only applies if you have deployed NSX Manager and CSM in Microsoft Azure using the Terraform scripts provided by NSX Cloud. See *Deploy NSX Cloud Components in Microsoft Azure using the NSX Cloud Marketplace Image* in the *NSX Installation Guide* for details on this deployment model.

Recovery Option for NSX Manager by Redeploying the NSX Manager image

NSX Manager appliance is deployed as a cluster of three NSX Manager nodes or NSX Manager VMs in your Microsoft Azure resource group. If you lose one of the NSX Manager nodes, you can simply redeploy that NSX Manager node using the image file available in your resource group. This option does not apply to recovering a lost CSM appliance because CSM is not deployed in a cluster and if lost, it cannot be recovered by redeploying because there are no other nodes replicating the data. See Redeploying NSX Manager from nsx_mgr_image in Microsoft Azure.

Recovery Options for NSX Manager and CSM Using the Microsoft Azure Recovery Vault Service

The Terraform scripts create a Microsoft Azure Recovery Vault service for backing up NSX Manager and CSM. The vault policy is named `<deployment_prefix>nsx-vault` and the default backup schedule is set to: `daily recurring at 11 p.m. UTC`. The backup policy triggers scheduled backup of all three NSX Manager nodes and also of the CSM appliance. You can edit the backup policy to suit your requirements.

See Restore CSM from Microsoft Azure Recovery Services Vault and Restore NSX Manager from Microsoft Azure Recovery Services Vault for details on recovery options.

Restore CSM from Microsoft Azure Recovery Services Vault

Microsoft Azure Recovery Services vault provides many different options for recovering VMs.

A few of these options are described in this section. Refer to Microsoft Azure documentation for more details.

- Restore OS Disk for CSM
- Restore to a new CSM VM retaining the private IP address when a new NIC must be provisioned

Restore OS Disk for CSM

You can use the Restore Disk option to recover CSM if you want to preserve other configurations such as the NIC.

1. In your Microsoft Azure subscription, navigate to the backup screen listing all available backups of the CSM VM, for example, **<deployment_prefix>-nsx-mgmt-rg > <deployment_prefix>-nsx-csm > From Operations in the left panel > Backup**.
2. Click **Restore VM** and select a Restore Point.
3. From **Restore Virtual Machine** select **Create new** under **Restore Configuration** and select **Restore disks** for **Restore type**.

   The NSX Cloud management VNet information is auto-populated.
4. Click **Restore**.

5 The restored OS disk is saved in the same resource group as the CSM VM. It is named as: **<deployment_prefix>-nsx-csm-<original OS disk name>-<backup timestamp>**, for example: `mynsxcsm-osdisk-20201022-191737`.

6 Swap the old CSM disk with the newly restored disk by navigating to **<deployment_prefix>-nsx-csm > Disks > Swap OS disk**.

Restore to a new CSM VM retaining the private IP address when a new NIC must be provisioned

If CSM NIC is not usable and must be re-provisioned, you must delete it, along with other CSM resources and restore CSM to a new VM by marking the private IP address as `static` in the deployment template. Follow these steps:

1 In your Microsoft Azure subscription, navigate to the backup screen listing all available backups of the CSM VM, for example, **<deployment_prefix>-nsx-mgmt-rg > <deployment_prefix>-nsx-csm > From Operations in the left panel > Backup**.

2 Click **Restore VM** and select a Restore Point.

3 From **Restore Virtual Machine** select **Create new** under **Restore Configuration** and select **Restore disks** for **Restore type**.

The NSX Cloud management VNet information is auto-populated.

4 Click **Restore**.

5 The restored OS disk is saved in the same resource group as the CSM VM. It is named as: **<deployment_prefix>-nsx-csm-<original OS disk name>-<backup timestamp>**, for example: `mynsxcsm-osdisk-20201022-191737`.

6 Delete the existing CSM VM and CSM NIC from the resource group.

a CSM VM

b CSM NIC and public IP

c CSM OS disk

d CSM data disk

7 Go to the OS disk restore created from the vault.

<deployment_prefix>-nsx-vault > Backup Items (under Protected Items in the left panel menu) > Azure Virtual Machine > <deployment_prefix>-nsx-csm > View all Jobs.

8 Click the CSM VM with the **Restore** operation. This is the OS disk recently restored. Click **Deploy Template**.

9 Update details as necessary for the new CSM VM.

10 Click **Edit Template**, change the private IP allocation method to `static` and provide the existing private IP address for the CSM appliance. For example,

```
"resources": [
     <...>,
     {
       <...>,
             "ipConfigurations": [
                 {
                     "properties": {
                         "privateIPAllocationMethod": "Static",
                         "privateIPAddress": "192.168.14.12"
                         <...>},
        <...>
```

11 Save the template, accept the terms and conditions, and click **Purchase**.

The new CSM VM is deployed, retaining the same IP addresses as the old CSM VM.

Restore NSX Manager from Microsoft Azure Recovery Services Vault

The Terraform script that deploys NSX Cloud components in your Microsoft Azure subscription, creates a Backup Policy using the Microsoft Azure Recovery Vault Services to back up all three NSX Manager nodes on a recurring basis.

If you want a backup of the only one node, you can edit the policy as necessary. The following instructions assume that all three NSX Manager nodes are backed up.

You have the following options for restoring the NSX Manager appliance deployed in a management VNet of your Microsoft Azure subscription:

- Restore OS Disk for NSX Manager
- Restore one NSX Manager node by creating a new NSX Manager VM
- Restore all three NSX Manager nodes by creating three new NSX Manager VMs

Restore OS Disk for NSX Manager

1 In your Microsoft Azure subscription, navigate to the backup screen listing all available backups of the NSX Manager VM, for example, **<deployment_prefix>-nsx-mgmt-rg > <deployment_prefix>-nsx-mgr0 > From Operations in the left panel > Backup**.

2 Click **Restore VM** and select a Restore Point.

3 From **Restore Virtual Machine** select **Create new** under **Restore Configuration** and select **Restore disks** for **Restore type**.

The NSX Cloud management VNet information is auto-populated.

4 Click **Restore**.

5 The restored OS disk is saved in the same resource group as the NSX Manager VM. It is named as: **<deployment_prefix>-nsx-mgr0-<original OS disk name>-<backup timestamp>**, for example: `mynsxmgr-osdisk-20201022-191737`.

6 Swap the old NSX Manager disk with the newly restored disk by navigating to **<deployment_prefix>-nsx-mgr > Disks > Swap OS disk**.

Restore one NSX Manager node by creating a new NSX Manager VM

You can restore this NSX Manager node to a new VM while retaining its private IP address. Restore the OS disk first so that you can create a template using the disk and then deploy a new VM that retains the NSX Manager's private IP address.

1 In your Microsoft Azure subscription, navigate to the backup screen listing all available backups of the NSX Manager VM, for example, if `nsx-mgr0` node is corrupted and you want to restore it from a back up, go to **<deployment_prefix>-nsx-mgmt-rg > <deployment_prefix>-nsx-mgr0 > From Operations in the left panel > Backup**.

2 Click **Restore VM** and select a Restore Point.

3 From **Restore Virtual Machine** select **Create new** under **Restore Configuration** and select **Restore disks** for **Restore type**.

The NSX Cloud management VNet information is auto-populated.

4 Click **Restore**.

5 The restored OS disk is saved in the same resource group as the NSX Manager VM. It is named as: **<deployment_prefix>-nsx-mgr>[0,1,2]-<original OS disk name>-<backup timestamp>**, for example: `mynsxmgr-osdisk-20201022-191737`.

6 Delete the existing NSX Manager entities from the resource group:

- NSX Manager VM
- NSX Manager NIC and public IP

  Note If you want to retain NSX Manager's public IP address, you can do that by reassigning the existing NSX Manager NIC to the new NSX Manager VM. Do not delete the NSX Manager NIC in that case. If NSX Manager NIC is not usable and you must delete it, then your new NSX Manager VM is assigned a new NIC and a new public IP address.

- NSX Manager OS disk
- NSX Manager data disk

7 Go to the OS disk restore created from the vault.

<deployment_prefix>-nsx-vault > Backup Items (under Protected Items in the left panel menu) > Azure Virtual Machine > <deployment_prefix>-nsx-mgr<[0,1,2]> > View all Jobs.

8 Click the CSM VM with the **Restore** operation. This is the OS disk recently restored. Click **Deploy Template**.

9 Update details as necessary for the new NSX Manager VM.

10 Click **Edit Template**, change the private IP allocation method to `static` and provide the existing private IP address for the NSX Manager appliance. For example,

```
"resources": [
     <...>,
     {
       <...>,
            "ipConfigurations": [
                {
                    "properties": {
                        "privateIPAllocationMethod": "Static",
                        "privateIPAddress": "192.168.43.11"
                        <...>},
       <...>
```

11 Save the template, accept the terms and conditions, and click **Purchase**.

The new NSX Manager VM is deployed, retaining the same IP address as the old NSX Manager VM.

12 Verify that the NSX Manager cluster is back to a healthy state. You can check cluster status using NSX CLI command `get cluster status` by logging in to any of the NSX Manager nodes using SSH. You can also check cluster status from the NSX Manager UI by going to **System > Appliances** in the NSX Manager UI.

Restore all three NSX Manager nodes

If all three of the NSX Manager nodes in the cluster must be restored, follow the same steps as described in Restore one NSX Manager node by creating a new NSX Manager VM to restore each node and then check cluster status.

Redeploying NSX Manager from `nsx_mgr_image` in Microsoft Azure

The Terraform script saves the images for NSX Manager and CSM in the Resource Group it creates in the NSX Cloud Management VNet.

You can use these images to redeploy NSX Manager or CSM.

This method of redeploying an image can help with recovering a lost or unusable NSX Manager node. However, you cannot use this method to recover a CSM. This is because NSX Manager is deployed in a three-node cluster and although CSM is joined with this cluster, CSM does not replicate NSX Manager data and NSX Manager nodes do not replicate CSM data. To recover CSM, follow the steps described in "Restore CSM from Microsoft Azure Recovery Services Vault" in the *NSX Administration Guide*.

Redeploying one NSX Manager node and Attaching it with the NSX Manager cluster

Use the following example for this procedure:

- You have the following NSX Manager nodes:
  - `Deployment1-NSX-MGR0`
  - `Deployment1-NSX-MGR1`
  - `Deployment1-NSX-MGR2`
- You lose NSX Manager node `Deployment1-NSX-MGR0`.

In the case that one NSX Manager node is lost, you can detach the defunct NSX Manager node, redeploy a new NSX Manager node using the image in your deployment's resource group, and then attach the newly deployed NSX Manager node to the NSX Manager cluster.

Follow these steps and refer to the example for identifying NSX Manager nodes:

1 To detach the defunct NSX Manager node from the NSX Manager cluster:

   a Log in to either of the working nodes over SSH and run the following NSX CLI command:

   ```
   Deployment1-NSX-MGR1> detach node <UUID of Deployment1-NSX-MGR0>
   ```

   b Check the status of the NSX Manager cluster; it shows stable with two healthy nodes:

   ```
   Deployment1-NSX-MGR1> get cluster status
   ```

2 To create a new NSX Manager node in your Microsoft Azure subscription:

   a Navigate to **Deployment1-nsx-mgmt-rg > Deployment1_nsx_mgr_image**.

   b Click **Create VM** and accept the pre-selected values for fields other than the ones specified in this table:

Parameter	Value
	Basic
Virtual machine name	Any descriptive name.
Size	The minimum requirement is: Standard_D4s_v3-4vcpus, 16 GB memory.
Authentication type	SSH
Username	Enter the default NSX Manager username: `nsxadmin`.
SSH Public Key Source	Select **Use existing public key** and copy-paste the public key for the NSX Manager node that you have detached from the cluster; from the example, for node `Deployment1-NSX-MGR0`.
	Disks
OS Disk type	Standard HDD

Parameter	Value
Data disks	Click **Create and attach a new disk** and select **Standard HDD.** for **Disk SKU** with a custom size of `100 GiB`. **Note** Ensure that the data disk host caching is set to read/write.
	Networking
Public IP	Click **Create new** and select **Static** for the **Assignment** option.
NIC network security group	Select **Advanced**
Configure network security group	Select the network security group created by the Terraform deployment for NSX Manager. From the example in this topic: `Deployment1-nsx-mgr-sg`
	Advanced

Custom data	Copy-paste the following, ensuring that you use your deployment's username and password:
	`#cloud-config`
	`hostname: ${hostname}`
	`bootcmd:`
	`- [cloud-init-per, instance, lvmdiskscan, lvmdiskscan]`
	`- [cloud-init-per, instance, secondary_partition, /opt/vmware/nsx-node-api/bin/set_secondary_partition.sh]`
	`chpasswd:`
	`expire: false`
	`list:`
	`- nsxadmin:<pwd>`
	`- root:<pwd>`

c Click **Review + create**.

The new NSX Manager node is deployed.

d Go to the newly deployed NSX Manager and set its private IP address setting to `static`.

3 Join the newly deployed NSX Manager with the existing NSX Manager cluster:

a Log in to the newly deployed NSX Manager node and run the following NSX CLI command to ensure it is up and running:

```
Deployment1-NSX-MGR0> get cluster status
```

b Join this NSX Manager with the cluster. You need the cluster id that you can retrieve from any of the other two NSX Manager nodes that are running:

```
Deployment1-NSX-MGR0> join <NSX-MGR0-IP> cluster-id <cluster-id> thumbprint <NSX-MGR0 api thumbprint> username <NSX-MGR0 username> password <NSX-MGR0 password>
```

c After the new NSX Manager node joins the cluster, run the command to check the status of the cluster with all three nodes:

```
Deployment1-NSX-MGR0> get cluster status
```

NSX Cloud FAQs and Troubleshooting

This topic covers some frequently asked questions and troubleshooting information.

How can I reinstall NSX Tools on Windows VM?

To reinstall NSX Tools on the Windows VM:

1 Uninstall the existing NSX Tools on the Windows VM. For details, see Uninstalling NSX Tools.

2 Reboot the Windows VM.

Important If you do not reboot the Windows VM after uninstalling the NSX Tools, reinstall can cause undesired behavior.

3 Reinstall the NSX Tools using the installation command. For details, see Install NSX Tools on Windows VMs .

How can I access the `nsxcli` commands after installing NSX Tools?

After installing NSX Tools on the Linux VM:

1 Log in to the Linux VM where you have installed NSX Tools.

2 Run the `sudo service nsx-agent-chroot nsx-exec bash` command. You will be directed to bash shell.

3 now run the `nsxcli` command. You will be directed to the `nsxcli` prompt.

You can now execute any required `nsxcli` commands like `get firewall rules` and so on.

After installing NSX Tools on the Windows VM:

1 Log in to the Windows VM where you have installed NSX Tools.

2 Open PowerShell.

3 On the PowerShell prompt, run the `nsxcli` command. You will be directed to the `nsxcli` prompt.

You can now execute any required `nsxcli` commands like `get firewall rules` and so on.

How can I verify that my NSX Cloud components are installed and running?

1 To verify that NSX Tools on your workload VM are connected to PCG, do the following:

   a Type the `nsxcli` command to open NSX CLI.

   b Type the following command to get the gateway connection status, for example:

```
get gateway connection status
Public Cloud Gateway  : nsx-gw.vmware.com:5555
Connection Status     : ESTABLISHED
```

2 The workload VMs must have the correct tags to connect to PCG:

   a Log in to the AWS console or the Microsoft Azure portal.

   b Verify the VM's eth0 or interface tag.

      The `nsx.network` key must have the value `default`.

My VMs launched using cloud-init are quarantined and do not allow installation of third-party tools. What should I do?

With the Quarantine Policy enabled, when launching VMs using cloud-init scripts with the following specifications, your VMs are quarantined upon launching and you are not able to install custom applications or tools on them:

- tagged with `nsx.network=default`
- custom services auto-installed or bootstrapped when the VM is powered on

Solution:

Update the `default` (AWS) or `default-vnet-<vnet-ID>-sg` (Microsoft Azure) security group to add inbound/outbound ports as required for the installation of custom or third-party applications.

I tagged my VM correctly and installed NSX Tools, but my VM is quarantined. What should I do?

If you encounter this problem, try the following:

- Check whether the NSX Cloud tag: `nsx.network` and its value: `default` are correctly typed in. This is case-sensitive.
- Resync the AWS or Microsoft Azure account from CSM:
  - Log in to CSM.
  - Go to **Clouds > AWS/Azure > Accounts**.
  - Click on **Actions** from the public cloud account tile and click **Resync Account**.

What should I do if I cannot access my workload VM?

From your Public Cloud (AWS or Microsoft Azure):

1 Ensure that all ports on the VM, including those managed by NSX Cloud, the OS firewall (Microsoft Windows or IPTables), and NSX are properly configured in order to allow traffic,

   For example, to allow `ping` to a VM, the following needs to be properly configured:

   - Security Group on AWS or Microsoft Azure. See Threat Detection using the NSX Cloud Quarantine Policy for more information.
   - NSX DFW rules. See Default Connectivity Strategy for NSX-Managed Workload VMs in the NSX Enforced Mode for details.

- Windows Firewall or IPTables on Linux.

2 Attempt resolving the issue by logging in to the VM using SSH or other methods, for example, the Serial Console in Microsoft Azure.

3 You can reboot the locked out VM.

4 If you still cannot access the VM, then attach a secondary NIC to the workload VM from which to access that workload VM.

Do I need a PCG even in the Native Cloud Enforced Mode?

Yes.

Can I change the IAM role for the PCG after I have onboarded my public cloud account in CSM?

Yes. You can rerun the NSX Cloud script applicable to your public cloud to regenerate the PCG role. Edit your public cloud account in CSM with the new role name after you regenerate the PCG role. Any new PCG instances deployed in your public cloud account will use the new role.

Note that existing PCG instances continue to use the old PCG role. If you want to update the IAM role for an existing PCG instance, go to your public cloud and manually change the role for that PCG instance.

Can I use the NSX on-premises licenses for NSX Cloud?

Yes, you can if your ELA has a clause for it.

I am using the URL from CSM to deploy PCG but I get an error because the gateway name is unresolvable.

When the URL from the CSM UI for installing PCG fails because of gateway name being unresolvable, do the following in the respective public cloud for the OS of your workload VM:

- On Microsoft Windows workload VMs in Microsoft Azure, run the following command and download the install script again using the URL from CSM:

```
Add-DnsClientNrptRule -Namespace "nsx-gw.vmware.local" -NameServers "168.63.129.16"
-DnsSecEnable
```

- On Microsoft Windows workload VMs in AWS, run the following command and download the install script again using the URL from CSM:

```
Add-DnsClientNrptRule -Namespace "nsx-gw.vmware.local" -NameServers "169.254.169.253"
-DnsSecEnable
```

- On Linux workload VMs in Microsoft Azure run the following command to get PCG's IP addresses and download the install script using these IP addresses with the URL from CSM.

```
nslookup nsx-gw.vmware.local 168.63.129.16 | awk '/^Address: / { print $2 }'
```

- On Linux workload VMs in AWS run the following command to get PCG's IP addresses and download the install script using these IP addresses with the URL from CSM.:

```
nslookup nsx-gw.vmware.local 169.254.169.253 | awk '/^Address: / { print $2 }'
```

How to connect CSM to MP using CA Certificate?

In the NSX Cloud setups, CSM connects to MP through a self-signed certificate. Instead of a self-signed certificate, you can use a CA-signed certificate, if required.

To use a CA-signed certificate, perform the following steps:

1 Log into the CSM appliance as a root user.

2 Copy the root `CA cert pem` file into CSM.

3 Get the Java KeyStore (JKS) password from the file as follows.

```
PASSWORD=`cat /config/http/.http_cert_pw`
```

4 Add the root CA certificate to the CSM JKS store using the following command.

```
keytool -importcert -file /root/myCA.pem -noprompt -alias nsx_mgmr_custom -storetype JKS
-keystore /usr/java/jre/lib/security/cacerts -storepass $PASSWORD
```

Note This example uses `/root/myCA.pem`. You must use path for your root `CA cert pem` file.

5 Check if alias is added using the following command.

```
keytool -list -v -keystore /usr/java/jre/lib/security/cacerts -storepass $PASSWORD | grep
        nsx_mgmr_custom
```

The command lists out the newly added CA certificates. This is used between CSM and NSX Manager.

The root CA certificate is now considered as valid, the CSM and NSX Manager can peer.