

ARISTA

Deployment Guide

Multi-domain Segmentation Services (MSS)

Version 1.2



Headquarters	Support	Sales
5453 Great America Parkway Santa Clara, CA 95054 USA +1-408-547-5500	+1-408-547-5502 +1-866-476-0000	+1-408-547-5501 +1-866-497-0000
www.arista.com/en/	support@arista.com	sales@arista.com

© Copyright 2025 Arista Networks, Inc. All rights reserved. The information contained herein is subject to change without notice. The trademarks, logos, and service marks ("Marks") displayed in this documentation are the property of Arista Networks in the United States and other countries. Use of the Marks is subject to the Arista Networks Terms of Use Policy, available at www.arista.com/en/terms-of-use. Use of marks belonging to other parties is for informational purposes only.


Contents

Chapter 1: Introduction.....	1
1.1 Definitions and Acronyms.....	4
Chapter 2: MSS Overview.....	5
Chapter 3: Initial Setup and Configuration.....	7
3.1 Getting Started.....	8
3.1.1 Requirements.....	9
3.1.2 Configuration.....	11
3.1.3 ZTX Configuration Examples.....	21
3.1.4 CloudVision Workflow.....	27
3.1.5 Onboard Devices.....	30
3.2 Installing the ZTX Monitor Node Appliance.....	35
3.2.1 Deploying the ZTX Monitor Node Workflow.....	36
3.2.2 ZTX Monitor Node Software Deployment.....	40
Chapter 4: Operational Aspects.....	48
4.1 MSS Dashboard.....	49
4.2 Policy Manager.....	53
4.3 Policy Monitor.....	60
4.4 Policy Builder.....	61
4.5 Policy Logs.....	68
4.6 MSS Studio.....	70
4.7 Create a Zero Trust Network with MSS.....	71
4.7.1 Prerequisites.....	72
4.7.2 Configuring Monitoring Rules.....	73
4.7.3 Configuring Security Policy Rules.....	75
4.7.4 Generating Coarse Rules.....	77
4.7.5 Generating Granular Rules.....	80
4.7.6 Implementing Zero Trust.....	83
Chapter 5: Troubleshooting MSS.....	84
5.1 Show Commands.....	85
5.2 Tracing.....	88
5.3 Considerations.....	89
5.3.1 ZTX-7250S-16S Monitor Node.....	89
5.3.2 Self-IP Support.....	89
5.3.3 Layer 2 Devices.....	89
5.3.4 Network Address Translation.....	89
5.3.5 Access Control Lists and Policy-Based Routing.....	90
Chapter 6: References.....	91
6.1 Related Documents.....	91

Introduction

Arista Multi-domain Segmentation Services (MSS) helps organizations overcome deficiencies in existing segmentation solutions and firmly establishes the network as the foundation of an effective zero-trust posture.

Figure 1-1: Four Components of MSS



Endpoint Identity and Microperimeter Tags

The first step in planning a microsegmentation strategy involves binding endpoints, workloads, and networks to specific microperimeter tags. MSS automates the management of microperimeters by connecting to external sources and using user-defined tags to assign endpoints into groups of endpoints requiring the same policy. MSS connects to external sources like NAC systems, CMDBs, and virtualization infrastructure management solutions like VMware vSphere.

"Zero Trust" Policy Planning with Traffic Map

Zero trust architecture principles require that all traffic on the network must be explicitly allowed by security policies. To create zero trust policies, it is vital to have complete visibility into existing traffic flows on the network. This complete visibility ensures that policies protect the right resources while at the same time not impeding legitimate business-justified flows. MSS maps all the communications within and across different parts of the network and provides a set of recommended policies to only permit trusted communications based on the observed traffic map.

Microperimeter Enforcement in the Network or Redirect to Firewall

MSS then distributes the zero trust policies to EOS-powered network switches. As traffic ingresses the switch, it can perform wire-speed distributed enforcement itself or redirect the traffic to any L3 firewall for stateful L4-7 inspection.

Importantly, Arista's switch-based enforcement overcomes the challenges associated with traditional ACL-based segmentation, such as TCAM exhaustion, by leveraging an advanced labeling engine that optimizes hardware utilization and maximizes scalability. Furthermore, because hardware labels are internal to a switch and are not shared across the network infrastructure, MSS can be seamlessly inserted into any multi-vendor network. This approach also avoids any proprietary protocols that force organizations into single-vendor networks.

MSS enforcement occurs at the network's edge, on the TOR switch, where traffic ingresses the MSS-enabled network.

Continuous Traffic Monitoring and Visibility of Policy Violations


After deploying the zero trust policies, MSS can monitor for policy violations, and report on the specific flows dropped in the network. This provides vital intelligence to the administrator to update the zero trust policies when they are valid. Yet, denying new services as well as monitor specific endpoints that are attempting to violate traffic rules.

ZTX Monitor Node Role

The primary purpose of a Multi-domain Segmentation Services (MSS) Monitor Node is to provide visibility into app-to-app traffic in the network, and to develop non-intrusive MSS policies that are aligned with applications requirements. It is a baseline capability for any micro or macro segmentation solution and is a practical way to build and deploy policies.

Configuring policies in terms of groups of end-points or prefixes is essential to simplify policy management and secure Data Center and Campus environments. The ZTX monitor node provides visibility into existing traffic needed to build such policies.

Figure 1-2: ZTX Node - Monitor Mode and Flow



Refer to [Deploying the ZTX Monitor Node](#) section for configuration information.

1.1 Definitions and Acronyms

- MSS - Multi-domain Segmentation Services
- CVP - CloudVision Portal
- CVaaS - CloudVision-as-a-Service
- TOR - Top of Rack
- STOR - Service Top of Rack
- ITOR - Ingress Top of Rack
- GRE - Generic Routing Encapsulation
- Tag - User-defined string of a property used in an end-point management system.
- Group or Field Set - Collection of end-points with the same tag or tags.
- Label - Hardware representation of a collection of tags internal to each switch.
- Security Domain - A security domain, a construct within MSS, manages the scope of a policy. For example, a security domain can be used for all devices within a geo-location, pod, or building.

MSS Overview


CloudVision provides support for microperimeter segmentation and enforcement as part of Arista's Multi-Domain Segmentation Service (MSS) for Zero Trust Networking (ZTN). ZTN works to reduce lateral movement into increasingly smaller areas where workloads are granularly identified and only approved connections are permitted.

MSS allows you to implement fine-grained security policies based on microperimeters defined around the identity of endpoints or applications. This enhances network security beyond what is possible when relying on traditional network boundaries like subnets and VRFs. Microperimeters minimize the overall attack surface available at any endpoint, providing a better defense against lateral attacks.

For guidance on using MSS to build a zero trust network, see [Create a Zero Trust Network with MSS](#).

Microperimeter segmentation and enforcement is possible through the use of CloudVision features like the MSS Studio, Policy Manager, Policy Builder, Policy Monitor, and MSS Dashboard. Use these features in the following way:

Figure 2-1: CloudVision Features



1. Use MSS Studio to define static groups and policy rules for a security domain or domain, create sensors, monitoring sessions, and security policies in CloudVision, and dynamically discover groups using the Policy Manager to define microperimeters. CloudVision discovers dynamic groups using various endpoint identity data sources, including Arista's network identity service (AGNI), VMware vCenter, configuration management databases (CMDBs) like ServiceNow, Infoblox, and CSV files. CloudVision learns about endpoints through these integrations and builds them into group member mappings stored in the CloudVision database. Configure traffic monitoring that allows CloudVision to generate security policy recommendations in the MSS Policy Builder and push the policy configuration to all devices in the domain VRF.
2. View suggested policies and configure security policies to view traffic sessions collected by the ZTX monitor node and define those security policies to explicitly allow only trusted traffic. CloudVision pushes the mirroring policies to Arista Network devices and Arista network devices mirrors the network traffic to Arista's ZTX appliance. The ZTX node performs stateful traffic analysis and sends session data back to CloudVision, where group member mappings are stored. Sessions are then mapped into group-to-group communication to produce a set of policy recommendations for users to review and audit.
3. Push the security policy configuration to network devices to distribute enforcement in the network or redirect to a firewall. The network devices enforce those security policies.
4. Continue to monitor traffic and update policies.

Important: CloudVision Portal (CVP) and CloudVision as a Service (CVaaS) support sensor creation.



- **CVP:** The built-in default sensor is sufficient to learn dynamic group information from all MSS data sources (e.g., AGNI, vCenter, ServiceNow, CSV, Infoblox). No additional standalone sensor is required.
- **CVaaS:** The default sensor can learn dynamic group information only from AGNI. A standalone sensor must be created and configured to support other MSS data sources (vCenter, ServiceNow, CSV, etc.).

See [Getting Started with MSS](#) for an end-to-end guide to the MSS workflow, from onboarding the ZTX monitor node to configuring and implementing network security policy rules.

Initial Setup and Configuration

This section is structured to provide a clear, step-by-step approach, ensuring a smooth and accurate initial setup. It covers essential procedures from the initial deployment to the integration of key components.

- [Getting Started](#)
 - [Requirements](#)
 - [Configuration](#)
 - [ZTX Configuration Examples](#)
 - [CloudVision Requirements](#)
 - [Onboard Devices](#)
- [Installing the MSS Monitor Node Appliance](#)
 - [Deploying the ZTX Monitor Node Workflow](#)
 - [ZTX Monitor Node Software Deployment](#)

3.1 Getting Started

MSS brings microperimeter segmentation and enforcement to CloudVision. Use new tools like the MSS Studio, Policy Manager, Policy Monitor, and Policy Builder to create groups used to define microperimeters and configure security policy rules that govern how traffic is forwarded among groups.

This end-to-end guide will help you get the most out of CloudVision's MSS functionality. You'll begin by onboarding EOS devices with MSS capability and the ZTX monitor node, then use the MSS Studio to define static groups and policy rules. Next, onboard data sources for dynamic group discovery and review and accept groups in the Policy Manager for use in policy building. Finally, configure a monitoring rule in the MSS Studio to enable you to generate and accept policy recommendations in the Policy Builder.

3.1.1 Requirements

To deploy MSS a combination of the following hardware is required:

Hardware

- TOR Switches
 - CCS-720XP, CCS-720D, CCS-722XPM
 - DCS-7010TX
 - DCS-7050X3
 - DCS-7280R3/R3A
- ZTX Appliances
 - ZTX-7250S-16S
 - vZTX Appliance



Note: vZTX is currently supported on KVM/ESXi VMs only.

Management and Visibility

- CVP or CVaaS for management and visibility.
 - CVP (with MSS Manager)



Note: CVP must be a three-node cluster.

Important: Enable Network Security - MSS in CVP and CVaaS by navigating to **Settings > Features** as shown in the following example.

Figure 3-1: CVP Settings - Features Network Security MSS



3.1.2 Configuration



Note: For CVP and EOS supported versions, please refer to the Arista MSS Datasheet - [Arista MSS Datasheet](#).

3.1.2.1 TOR Configuration

Perform the following configuration steps to prepare the devices for MSS using the CLI. Several configuration steps are required on all supported platforms, while others are unique to specific SKUs.

Refer to the specific [Strata](#) and [Sand](#) sections for SKU-specific configuration information.



Note: These common and unique configurations on Sand and Strata devices can be added manually via CLI and reconciled to CVP or added as a configlet in CVP and pushed to the devices.

Common Configurations for Sand and Strata

For the devices to be able to stream to CVP and receive information from CVP, enable gRPC Network Management Interface (gNMI) on the devices:

```
ld459.12:57:58(config)# management api gnmi
ld459.12:58:13(config-mgmt-api-gnmi)# transport grpc default
ld459.12:58:20(config-gnmi-transport-default)# no shutdown
```

If the management is in a user-defined **vrf**, for example **vrf mgmt** tenable gNMI in that user-defined vrf.

```
ld459.12:59:16(config)# management api gnmi
ld459.12:59:18(config-mgmt-api-gnmi)# transport grpc default
ld459.12:59:25(config-gnmi-transport-default)# vrf mgmt
ld459.12:59:40(config-gnmi-transport-default)# no shutdown
```

For devices to receive dynamic group information from CVP, enter the following parameters at the end of the **TerminAttr** configuration:

```
ld459.13:00:28(config)# daemon TerminAttr
ld459.13:00:29(config-daemon-TerminAttr)# exec /usr/bin/TerminAttr -smashecludes=ale,flexCou
nter,hardware,
kni,pulse,strata -cvaddr=10.248.18.240:9910,10.248.18.241:9910,10.248.18.242:9910 -cvauth=token,/
tmp/token
-cvvrf=default -taillogs -cvtargetconfigs arista/traffic-policy
! New daemon configuration will only take effect after restarting the daemon by doing 'shutdown/no
shutdown'.
ld459.13:00:39(config-daemon-TerminAttr)# no shutdown
```

3.1.2.1.1 Traffic Policy Configuration

As part of enforcing MSS policies, CloudVision configures the traffic policy rules for individual EOS devices.

If required, refer to the [Arista EOS System Configuration Guide](#) for more information on traffic policies.

3.1.2.1.2 Strata Configuration

Perform the following configuration steps to prepare the devices for MSS using the CLI. Several unique configuration steps are required on the Strata platform. These apply to:

- DCS-7050X3
- CCS-720DP (excluding 720DP-24S)
- CCD-720DT (excluding 720DT-24S)
- CCS-720DF
- CCS-720XP (excluding 720XP-96ZC2, 720XP-48TXH-2C-S)
- CCS-722XPM (excluding 722XPM-48ZY8)
- DCS-7010TX

Refer to the specific [Strata](#) and [Sand](#) sections for detailed information.

Note: These common and unique configurations on Sand and Strata devices can be added manually via CLI and reconciled to CVP or added as a configlet in CVP and pushed to the devices.

On Strata devices (7050S, 720S, and 722S platforms) add the following configuration to enable MSS:



```
ld459.13:02:20(config)# traffic-policies
ld459.13:04:57(config-traffic-policies)# transforms interface prefix common source-destination
```



Attention: After adding transforms via the CLI you must reboot the device.



Important: Enable IP routing on Strata devices when using MSS.

3.1.2.1.3 Sand Configuration

Perform the following configuration steps to prepare the devices for MSS using the CLI. Several unique configuration steps are required on the Sand platform. These apply to:

- DCS-7280R3/R3A (except 7280SR3-40YC6, 7280SR3E-40YC6 and 7280TR3-40C6)

Refer to the specific [Strata](#) and [Sand](#) sections for detailed information.



Note: These common and unique configurations on Sand and Strata devices can be manually added via CLI, reconciled to CVP, or added as a configlet in CVP and pushed to the devices.

On Sand devices (7280R3 platforms), use the [MSS Base Profile](#). The base profile supports:

- **L4-Port Transforms:** Allows the configuration of L4-port-based match rules.
- **Self-IP Support:** Allows the configuration of a self-IP-based rule if the policy contains **deny any any** at the end to save Control Plane packets from hitting this **deny any any** rule.
- **Mirror Action:** Allows the creation of Monitor sessions to achieve ZTX monitoring node functionality.

3.1.2.1.3.1 Configuring the TCAM profile

TCAM profiles should follow a specific flow and order, addressing required features, key fields, and important actions, as follows:

Required Features

1. `feature traffic-policy port ipv4`
2. `feature traffic-policy port ipv6`
3. `feature traffic-policy vlan-interface ipv4`
4. `feature traffic-policy vlan-interface ipv6`

Important Key Fields

1. `dst-port`
2. `dst-ip-label`
3. `src-ip-label`
4. `ip-protocol`
5. `l4-dst-port-label`
6. `l4-src-port-label`
7. `vxlan-decapsulated`

Important Actions

1. `count`
2. `drop`
3. `mirror`
4. `redirect`



Important: Based on the specific requirements, removing some actions and adding others may be necessary. For example, to have the action `log` (not MSS use case), remove the `redirect/mirror` action and add the `log` action.

The `feature traffic-policy port ipv4` and `feature traffic-policy port ipv6` commands define the IPv4 and IPv6 traffic-policy lookups, respectively. In the later example, the key fields include **`dst-ip-label`**, **`src-ip-label`**, **`dst-ipv6-label`**, and **`src-ipv6-label`**. This means the hardware is programmed to first do a lookup on the **`dst-ip`**, **`src-ip`**, etc., to generate a label that is then used in the rules TCAM lookup.

The example shown in [TCAM Profile Supported Systems](#) includes **I4-src-port** and **I4-dst-port** in the key field, which means the I4 source and destination port are directly used in the rules TCAM lookup.

Alternatively, as shown in the following snippet, if the key field contains **I4-src-port-label** and **I4-dst-port-label** instead of **I4-src-port** and **I4-dst-port**, this would mean there is a hardware lookup to transform the I4 source and destination ports into labels that are then used in the rules TCAM lookup.

Since the I4 source and destination port lookups consume a TCAM bank each, changing a profile will result in using two extra banks for the port transform tables. However, depending on the actual traffic policy, this might result in lower TCAM utilization for the rules TCAM lookup. The system must analyze the actual utilization of the TCAM banks for a given TCAM profile and its traffic-policy configuration before choosing the TCAM profile configuration.

TCAM profiles are supported on 7280R3 and 7500R3. If you are using another platform, skip this section.

```
lyv571.13:02:41 (config) # hardware tcam
lyv571.13:02:48 (config-tcam) # profile mss-traffic-policy-l4-xform
lyv571.13:02:51 (config-tcam-profile-mss-traffic-policy-l4-xform) #
```

Configure the TCAM profile as required. For MSS, configure the TCAM features using the [MSS Base Profile](#).

3.1.2.1.3.2 MSS Base Profile

TCAM Profile on 7280R3 platforms.

```

hardware tcam
  profile mss-traffic-policy-l4-xform
    feature acl port mac
      key field dst-mac ether-type src-mac
      action count drop mirror
      packet ipv4 forwarding bridged
      packet ipv4 forwarding routed
      packet ipv4 forwarding routed multicast
      packet ipv4 mpls ipv4 forwarding mpls decap
      packet ipv4 mpls ipv6 forwarding mpls decap
      packet ipv4 non-vxlan forwarding routed decap
      packet ipv4 vxlan forwarding bridged decap
      packet ipv6 forwarding bridged
      packet ipv6 forwarding routed
      packet ipv6 forwarding routed decap
      packet ipv6 forwarding routed multicast
      packet ipv6 ipv6 forwarding routed decap
      packet mpls forwarding bridged decap
      packet mpls ipv4 forwarding mpls
      packet mpls ipv6 forwarding mpls
      packet mpls non-ip forwarding mpls
      packet non-ip forwarding bridged
      sequence 55
      key size limit 160
    feature forwarding-destination mpls
      sequence 100
    feature mirror ip
      key field dscp dst-ip ip-frag ip-protocol l4-dst-port l4-ops l4-src-port src-ip tcp-
control
      action count mirror set-policer
      packet ipv4 forwarding bridged
      packet ipv4 forwarding routed
      packet ipv4 forwarding routed multicast
      packet ipv4 non-vxlan forwarding routed decap
      sequence 80
      key size limit 160
    feature mpls
      action drop redirect set-ecn
      packet ipv4 mpls ipv4 forwarding mpls decap
      packet ipv4 mpls ipv6 forwarding mpls decap
      packet mpls ipv4 forwarding mpls
      packet mpls ipv6 forwarding mpls
      packet mpls non-ip forwarding mpls
      sequence 5
      key size limit 160
    feature mpls pop ingress
      sequence 90
    feature qos ip
      key field dscp dst-ip ip-frag ip-protocol l4-dst-port l4-ops l4-src-port src-ip tcp-
control
      action set-dscp set-policer set-tc
      packet ipv4 forwarding routed
      packet ipv4 forwarding routed multicast
      packet ipv4 mpls ipv4 forwarding mpls decap
      packet ipv4 mpls ipv6 forwarding mpls decap
      packet ipv4 non-vxlan forwarding routed decap
      sequence 75
      key size limit 160
    feature qos ipv6
      key field dst-ipv6 ipv6-next-header ipv6-traffic-class l4-dst-port l4-src-port src-ipv6-
high src-ipv6-low
      action set-dscp set-policer set-tc
      packet ipv6 forwarding routed
      sequence 70
    feature traffic-policy port ipv4
      key field dst-port dscp dst-ip-label ip-frag ip-fragment-offset ip-length ip-protocol l4-
dst-port-label l4-src-port-label src-ip-label tcp-control ttl vxlan-decapsulated
      action count drop mirror redirect
      packet ipv4 forwarding bridged
      packet ipv4 forwarding routed
      packet ipv4 mpls ipv4 forwarding mpls decap
      packet ipv4 non-vxlan forwarding routed decap
      packet mpls ipv4 forwarding bridged
      packet mpls ipv4 forwarding mpls

```

```

sequence 45
key size limit 160
feature traffic-policy port ipv4 egress
key field dscp dst-ip-label ip-frag ip-protocol l4-dst-port l4-src-port src-ip-label tcp-
control
action count drop log
packet ipv4 forwarding routed
packet mpls ipv4 forwarding mpls
key size limit 160
feature traffic-policy port ipv6
key field dst-port dst-ipv6-label hop-limit ipv6-length ipv6-next-header ipv6-traffic-
class l4-dst-port-label l4-src-port-label src-ipv6-label tcp-control
action count drop mirror redirect
packet ipv4 mpls ipv6 forwarding mpls decap
packet ipv6 forwarding bridged
packet ipv6 forwarding routed
packet ipv6 forwarding routed decap
packet mpls ipv6 forwarding bridged
packet mpls ipv6 forwarding mpls
sequence 25
feature traffic-policy port ipv6 egress
key field dscp dst-ipv6-label ipv6-next-header l4-dst-port l4-src-port src-ipv6-label
tcp-control
action count drop log
packet ipv6 forwarding routed
packet mpls ipv6 forwarding mpls
feature traffic-policy vlan-interface ipv4
key field dst-port dscp dst-ip-label ip-frag ip-fragment-offset ip-length ip-protocol l4-
dst-port-label l4-src-port-label src-ip-label tcp-control ttl vxlan-decapsulated
action count drop mirror redirect
packet ipv4 forwarding routed
packet ipv4 forwarding bridged
key size limit 160
feature traffic-policy vlan-interface ipv6
key field dst-port dst-ipv6-label hop-limit ipv6-length ipv6-next-header ipv6-traffic-
class l4-dst-port-label l4-src-port-label src-ipv6-label tcp-control
action count drop mirror redirect
packet ipv6 forwarding routed
packet ipv6 forwarding bridged
key size limit 160
feature tunnel vxlan
packet ipv4 vxlan eth ipv4 forwarding routed decap
packet ipv4 vxlan forwarding bridged decap
sequence 50
key size limit 160

```



Note: The TCAM Base Profile supports mirror and redirect actions.

3.1.2.1.3.3 TCAM Profile Supported Systems

TCAM Profile on 7500R3 and 7800R3 systems.

The following is an example of a custom profile with interface traffic policy supported:

```

hardware tcam
  profile traffic-policy
    feature acl port mac
      sequence 55
      key size limit 160
      key field dst-mac ether-type src-mac
      action count drop
      packet ipv4 forwarding bridged
      packet ipv4 forwarding routed
      packet ipv4 forwarding routed multicast
      packet ipv4 mpls ipv4 forwarding mpls decap
      packet ipv4 mpls ipv6 forwarding mpls decap
      packet ipv4 non-vxlan forwarding routed decap
      packet ipv4 vxlan forwarding bridged decap
      packet ipv6 forwarding bridged
      packet ipv6 forwarding routed
      packet ipv6 forwarding routed decap
      packet ipv6 forwarding routed multicast
      packet ipv6 ipv6 forwarding routed decap
      packet mpls forwarding bridged decap
      packet mpls ipv4 forwarding mpls
      packet mpls ipv6 forwarding mpls
      packet mpls non-ip forwarding mpls
      packet non-ip forwarding bridged
    feature forwarding-destination mpls
      sequence 100
    feature mirror ip
      sequence 80
      key size limit 160
      key field dscp dst-ip ip-frag ip-protocol l4-dst-port l4-ops l4-src-port src-ip tcp-
control
      action count mirror set-policer
      packet ipv4 forwarding bridged
      packet ipv4 forwarding routed
      packet ipv4 forwarding routed multicast
      packet ipv4 non-vxlan forwarding routed decap
      feature mpls
      sequence 5
      key size limit 160
      action drop redirect set-ecn
      packet ipv4 mpls ipv4 forwarding mpls decap
      packet ipv4 mpls ipv6 forwarding mpls decap
      packet mpls ipv4 forwarding mpls
      packet mpls ipv6 forwarding mpls
      packet mpls non-ip forwarding mpls
      feature pbr ip
      sequence 60
      key size limit 160
      key field dscp dst-ip ip-frag ip-protocol l4-dst-port l4-ops-18b l4-src-port src-ip tcp-
control
      action count redirect
      packet ipv4 forwarding routed
      packet ipv4 mpls ipv4 forwarding mpls decap
      packet ipv4 mpls ipv6 forwarding mpls decap
      packet ipv4 non-vxlan forwarding routed decap
      packet ipv4 vxlan forwarding bridged decap
      feature pbr ipv6
      sequence 30
      key field dst-ipv6 ipv6-next-header l4-dst-port l4-src-port src-ipv6-high src-ipv6-low
tcp-control
      action count redirect
      packet ipv6 forwarding routed
      feature pbr mpls
      sequence 65
      key size limit 160
      key field mpls-inner-ip-tos
      action count drop redirect
      packet mpls ipv4 forwarding mpls
      packet mpls ipv6 forwarding mpls
      packet mpls non-ip forwarding mpls

```

```

feature qos ip
sequence 75
key size limit 160
key field dscp dst-ip ip-frag ip-protocol l4-dst-port l4-ops l4-src-port src-ip tcp-
control
action set-dscp set-policer set-tc
packet ipv4 forwarding routed
packet ipv4 forwarding routed multicast
packet ipv4 mpls ipv4 forwarding mpls decap
packet ipv4 mpls ipv6 forwarding mpls decap
packet ipv4 non-vxlan forwarding routed decap
feature qos ipv6
sequence 70
key field dst-ipv6 ipv6-next-header ipv6-traffic-class l4-dst-port l4-src-port src-ipv6-
high src-ipv6-low
action set-dscp set-policer set-tc
packet ipv6 forwarding routed
feature traffic-policy port ipv4
sequence 45
key size limit 160
key field dscp dst-ip-label icmp-type-code ip-frag ip-fragment-offset ip-length ip-
protocol l4-dst-port l4-src-port src-ip-label tcp-control ttl
action count drop log set-dscp set-tc
packet ipv4 forwarding routed
feature traffic-policy port ipv6
sequence 25
key field dst-ipv6-label hop-limit icmp-type-code ipv6-length ipv6-next-header ipv6-
traffic-class l4-dst-port l4-src-port src-ipv6-label tcp-control
action count drop log set-dscp set-tc
packet ipv6 forwarding routed
feature tunnel vxlan
sequence 50
key size limit 160
packet ipv4 vxlan eth ipv4 forwarding routed decap
packet ipv4 vxlan forwarding bridged decap

```

3.1.3 ZTX Configuration Examples

This chapter provides several configuration examples that can be used to implement both out-of-band and in-band communication requirements for ZTX.

Out-of-band Configuration Prerequisites

As a prerequisite, the ZTX requires an IP address assigned to its management interface and an active Streaming Agent (**TerminAttr**) instance, in order to communicate with CloudVision. Refer to the documentation to understand the different options to onboard an Arista device in CloudVision.

This example assumes that:

- The CloudVision cluster uses the following IP addresses: 172.28.137.75, 172.28.130.47, 172.28.133.90
- The ZTX device is provisioned with at least one valid NTP server and with the proper clock time-zone, for example:

```
ntp server my-ntp-server.mydomain.mycompany.com
!  
clock timezone US/Pacific
!
```

- The ZTX device is provisioned with a management address of 172.28.137.229/20 and a default gateway in the default VRF:

```
interface Management1/1  
ip address 172.28.137.229/20  
!  
ip route 0.0.0.0/0 172.28.128.1
```




Note: The existing code referenced earlier reflects the IP addresses of CloudVision and the VRF value (default) of the management interface.

In-band Configuration - MLAG Peering and Static Routing

This example is based on the following topology diagram, where the ZTX is physically adjacent to a pair of Arista switches configured as MLAG peers, and a single port channel group is established between them.

Figure 3-2: Example - Topology LAG Peering and Static Routing



These are in summary the configuration steps for this example:

1. A single port channel is provisioned on both the ZTX and MLAG pair.
2. The ZTX is configured with a unique IP address in a specific subnet that is assigned to an SVI.
3. The VLAN used by this SVI is configured on the port channel on both the ZTX and the MLAG pair.
4. The same IP subnet is assigned to the corresponding SVI on the MLAG pair.
5. The ZTX is configured with a static route in order to have reachability to the switches in the Security Domain.
6. The MLAG pair communicates with the rest of the IP network with a dynamic protocol of choice and advertises the local subnet of the SVI to provide reachability to the ZTX IP address.
7. The ARP timeout of the ZTX is tuned to achieve peer adjacency persistence.

The following are the configuration step details:

1. A single port channel is provisioned on both the ZTX and MLAG pair.

On ZTX:

```
interface Port-Channel16
!
interface Ethernet1/1 - Ethernet1/16
channel-group 16 mode active
!
```

On MLAG left and right switches:

```
interface Port-Channel8
mlag 8
!
interface Ethernet11/1 - 4
speed forced 10000full
channel-group 8 mode active
!
interface Ethernet13/1 - 4
speed forced 10000full
channel-group 8 mode active
!
```

2. The ZTX is configured with a unique IP address in a specific subnet that is assigned to an SVI.

```
vlan 1016
!
interface vlan1016
ip address 10.10.16.4/29
!
```

3. The VLAN used by this SVI is configured on the port channel on both the ZTX and the MLAG pair.

On the ZTX:

```
interface Port-Channel16
switchport access vlan 1016
!
```

On the MLAG switch pair:

```
vlan 1016
!
interface Port-Channel8
switchport access vlan 1016
!
```

4. The same IP subnet is assigned to the corresponding SVI on the MLAG pair.

On both left and right switch:

```
interface vlan1016
ip virtual address 10.10.16.1/29
!
```

5. The ZTX is configured with a static route in order to have reachability to the switches in the Security Domain.

Assuming these switches use addresses taken from a 10.10.0.0/19 aggregate subnet:

```
ip routing
!  
ip route 10.0.0.0/19 10.10.16.1
```

6. The MLAG pair communicates with the rest of the IP network with a dynamic protocol of choice and advertises the locally connected subnet of the SVI to provide reachability to the ZTX IP address.

For example, using OSPF:

```
router ospf 10  
redistribute connected  
!
```


7. The ARP timeout of the ZTX is tuned to achieve peer adjacency persistence This step is recommended with static routing and ensures that the layer-2 adjacency does not expire in case monitoring is inactive and peer links are idle.

```
arp aging timeout default 180
```

In-band Configuration - Dynamic Routing

This example is based on the following topology diagram, where the ZTX is physically adjacent to two or more Arista switches, and one channel group per switch is established between them. In case the peering switches are two, they can optionally form an MLAG pair.

Figure 3-3: Example - Topology of Layer-3 Peering and Dynamic Routing



These are in summary the configuration steps for this example:

1. A single routed port channel is provisioned on the ZTX and each peering switch, and configured with a point-to-point subnet.
2. The ZTX is configured with a unique host IP address that is assigned to a loopback interface.
3. The ZTX is configured with a dynamic routing protocol and peering with the adjacent switches, in order to have mutual reachability between its loopback address and those assigned to the switches in the Security Domain.

The following are the configuration step details:

1. A single routed port channel is provisioned on the ZTX and each peering switch, and configured with a point-to-point subnet.

```
interface Port-Channel8
no switchport
ip address 192.168.100.101/31
!
interface Ethernet1/1 - Ethernet1/8
channel-group 8 mode active
!
interface Port-Channel16
no switchport
ip address 192.168.100.103/31
!
interface Ethernet1/9 - Ethernet1/16
channel-group 16 mode active
!
```

On left peering switch:

```
interface Port-Channel8
no switchport
ip address 192.168.100.100/31
!
interface Ethernet11/1 - 4
speed forced 10000full
channel-group 8 mode active
!
interface Ethernet13/1 - 4
speed forced 10000full
channel-group 8 mode active
!
```

On right peering switch:

```
interface Port-Channel16
no switchport
ip address 192.168.100.102/31
!
interface Ethernet11/1 - 4
speed forced 10000full
channel-group 16 mode active
!
interface Ethernet13/1 - 4
speed forced 10000full
channel-group 16 mode active
!
```

2. The ZTX is configured with a unique host IP address that is assigned to a loopback interface.

```
interface Loopback0
description router-id
ip address 10.135.2.16/32
!
```

3. The ZTX is configured with a dynamic routing protocol, in this example BGP, and peering with the adjacent switches, in order to have mutual reachability between its loopback address and those assigned to the switches in the Security Domain.

On ZTX:

```
router bgp 64516
router-id 10.135.2.16
distance bgp 20 200 200
maximum-paths 2
neighbor UNDERLAY peer group
neighbor UNDERLAY maximum-routes 120
neighbor 192.168.100.100 peer group UNDERLAY
neighbor 192.168.100.100 remote-as 64504
neighbor 192.168.100.100 description SwitchLeft
neighbor 192.168.100.102 peer group UNDERLAY
neighbor 192.168.100.102 remote-as 64504
neighbor 192.168.100.102 description SwitchRight
redistribute connected
!
address-family ipv4
neighbor UNDERLAY activate
!
```

On left peering switch:

```
router bgp 64504
```

```

network 10.135.2.0/24
neighbor ZTX peer group
neighbor ZTX route-map LOOPBACKS out
neighbor 192.168.100.101 peer group ZTX
neighbor 192.168.100.101 remote-as 64516
neighbor 192.168.100.101 description ZTX-1
!
address-family ipv4
neighbor ZTX activate
!
route-map LOOPBACKS permit 10
match ip address prefix-list LOOPBACKS
!
ip prefix-list LOOPBACKS seq 5 permit 10.135.2.0/24
!

```

On right peering switch:

```

router bgp 64504
network 10.135.2.0/24
neighbor ZTX peer group
neighbor ZTX peer group
neighbor ZTX route-map LOOPBACKS out
neighbor 192.168.100.103 peer group ZTX
neighbor 192.168.100.103 remote-as 64516
neighbor 192.168.100.103 description ZTX-1
!
address-family ipv4
neighbor ZTX activate
!
route-map LOOPBACKS permit 10
match ip address prefix-list LOOPBACKS
!
ip prefix-list LOOPBACKS seq 5 permit 10.135.2.0/24
!

```

3.1.4 CloudVision Workflow

CloudVision Portal (CVP) is your turnkey solution for managing, monitoring, and maintaining your network devices. Once you've onboarded your switches in CVP review the following steps and procedures.



Note: The following workflow also applies when using CloudVision-as-a-Service (CVaaS).

Figure 3-4: Workflow



Use following the QR Code or the links to locate the relevant CVP sections providing detailed configuration information.

Figure 3-5: CVP Documentation QR Code



CVP Documentation Link - <https://www.arista.io/help/articles/b3ZlcnZpZXcubXNzLmdldHRpbmdTdGFydGVk>

Security Domains - Under **Security Domains**, select **Add Security Domain** and enter a title for the domain.

[Onboarding a Data Source](#) - A data source is a third-party device or management system that streams data to CloudVision. The required configuration is managed in CloudVision.

Create Policies - Static and Dynamic - **Under Policies**, select **Add Policy** and enter a policy name and then apply the policy on the security domain VRF.

3.1.5 Onboard Devices

Onboard 7280R3 and 7050X3/720/722-series EOS devices with MSS capability and the [ZTX 7250S monitor appliance](#) and register them for use in Studios.

1. [Onboard](#) 7280R3 and 7050X3/720/722-series devices and the ZTX monitor appliance.



Note: If the End-to-End Provisioning toggle is enabled in [General Settings](#), newly-onboarded devices will automatically be available to be registered for use in Studios.

2. Register devices in Studios

Once you've onboarded relevant EOS devices, you'll register them in the Inventory and Topology Studio for use in Studios.

- Navigate to the [Inventory and Topology Studio](#) and click **Network Updates**.
- Enable the checkbox next to the onboarded device or devices and click **Accept Updates**.

You'll now be able to use these devices in the [MSS Studio](#) where you'll configure static security domain groups and security policy rules.


3.1.5.1 Reconciling Device Configuration

Reconcile is used to resolve differences between device designed configuration and the device running configuration. This is typically required when a device's running config has been updated via CLI instead of CloudVision.

Building the workspace will detect if any device configuration is non-compliant and needs to be reconciled.

1. Select **Reconcile**.

Figure 3-6: Reconcile Tab



Devices detected with non-compliant running configuration will be displayed.


Figure 3-7: Start Build



This will update the list of devices that are out of compliance.

3. Select one or more devices.

Figure 3-8: Select Devices




Choose the lines from the running configuration to include or overwrite the designed configuration using the **checkboxes**.



Tip: Select **Auto-Reconcile** to reconcile all devices with one action.


4. Select **Reconcile**.

Figure 3-9: Reconcile Multiple Devices



The workspace is rebuilt, and a new container in the Container Tree called **Reconciled Configlets** is displayed.

Figure 3-10: Configlets



The order of configuration is important. Ensure that the reconciled configuration follows the established EOS hierarchy both within the configlet and in the order of containers.

5. Select **Review Workspace** and submit the workspace to change control.

When the associated change control is executed, the running configuration and the designed configuration will be synchronized, bringing the running configuration of the device into compliance.

3.2 Installing the ZTX Monitor Node Appliance

For Installation and Configuration of the MSS Monitor Node appliance (ZTX-72XXS), refer to the [Quick Start Guide](#) on the Arista documentation page.

3.2.1 Deploying the ZTX Monitor Node Workflow

On deploying a ZTX monitor node and following the user workflow to trigger mirroring traffic to the node, the solution does the following:

1. CVP configures mirroring: CVP installs MSS rules on the ingress TORs to mirror traffic from selected network VRFs or network devices towards the ZTX node.
2. ITORs mirror traffic to the ZTX node: Once the mirroring configuration is active, selected traffic is mirrored from the ingress TOR via a GRE tunnel to the ZTX node.
3. ZTX node summarizes session: Upon receiving mirrored traffic from the ingress TORs, the ZTX node tracks and aggregates sessions bidirectionally.
4. ZTX node exports summarized sessions to CVP: The ZTX node exports the summarized session information to CVP
5. CVP associates summarized session data with meta data to suggest policy that the user can then review and modify.
6. CVP suggests a permit action to create MSS rules based on these summarized session records. However, upon review, the policy action can be modified and pushed to the ITOR switches, enforcing the rules.

3.2.1.1 Configuration

The ZTX appliance is managed primarily through CloudVision's **MSS Service Studio** once the device has been onboarded and bootstrapped with configuration pushed through the **Static Configuration Studio**. Device onboarding is described by the following [Device Onboarding](#) instructions. Since **MSS Service** will manage the device at runtime, this configuration section will primarily focus on the necessary bootstrap configuration that is pushed through **Static Configlets Studio**. Configuration pushed through **MSS Service** can also be validated.

3.2.1.2 Static Configuration Studio

Static Configuration Studio can be leveraged to build the initial configurations on the ZTX node and include provisioning loopback, ethernet and port-channel interface and IP connectivity to other devices in the network. Additional details about Static Configlet Studios are available on the Arista CloudVision Help Center [Static Configuration Studio](#) site.

Ports from the ZTX device used as uplinks to the MLAG Service TOR should all be members of a single port-channel to achieve maximum throughput with the service TOR. A loopback interface is needed to act as the GRE tunnel endpoint to terminate mirrored packets. The following is a sample configuration that can be pushed from **Static Configuration Studio** to bootstrap the device so that it can be managed by **MSS Service**.

```
ZTX# show running-config
!
! Port-Channel with member ports connected to MLAG Service TOR.
! MLAG service TOR peers should have an IP address in the same network
interface Port-Channell
no switchport
ip address 10.10.250.1/24
!
interface Ethernet1/1
speed forced 10000full
no switchport
channel-group 1 mode active
!
interface Ethernet1/2
speed forced 10000full
no switchport
channel-group 1 mode active
!
! Loopback Interface used as GRE tunnel endpoint.
interface Loopback0
ip address 10.10.254.1/32
!
! Default static route for reachability of GRE mirror source
ip route 0.0.0.0/0 10.10.250.2
```

3.2.1.3 CloudVision MSS Service

Use the **MSS Service** to configure the ZTX device and the mirroring session on the MSS TOR. The following section describes the configuration that **MSS Service** will push and how the configurations contribute to the overall solution, aside from show commands for debugging the user is not expected to actively manage the ZTX device after the onboarding and bootstrap stage are complete.

The following configuration puts the ZTX device into monitor mode.

```
ZTX# show running-config
!
firewall distributed instance
mode monitor
no disabled
!
```

The following configuration exports the flows observed by the monitoring device to CV using IPFIX as the transport. Note that TA is acting as IPFIX collector listening on 127.0.0.1 for IPFIX records and then streaming that data to CV for ingest to provide telemetry and policy recommendations.

```
ZTX# show running-config
!
flow tracking firewall distributed
tracker flowtrkr
exporter exp
collector 127.0.0.1
local interface Loopback0
no shutdown
!
```

Tunnels are configured to terminate each GRE tunnel created from a TOR's monitor sessions, and flow tracker is enabled on all the tunnels. The tunnel source is always set to the IP address of the loopback interface created during the bootstrap process. The tunnel destination corresponds to the GRE tunnel endpoint of the TOR that originated the mirror session. Note that mirrored traffic terminated at the ZTX device and is always dropped so the reverse GRE tunnel is never utilized.

```
ZTX# show running-config
!
interface Tunnel0
flow tracker firewall distributed flowtrkr
tunnel mode gre
tunnel source 10.10.254.1
tunnel destination 10.10.254.2
!
```



Note: Configurations of tunnel interface and flow tracker come from CVP MSS Service. However, configurations for interfaces connecting to STOR and Loopback Interface should be done via CLI/Static Configlets Studio as a prerequisite to using MSS Service.

3.2.2 ZTX Monitor Node Software Deployment

3.2.2.1 Deploying the ZTX Monitor Node


Deploying the ZTX Monitor Node in an Arista Network

This chapter describes how the ZTX can be physically and logically inserted in an Arista network and provides reference configuration examples.

IP Communication Requirements

As represented in the following logical diagram, there are two fundamental IP communication requirements for the ZTX Traffic Mapper:

Figure 3-11: ZTX IP Communication Requirements



1. The ZTX requires bidirectional IP communication with Arista Cloud Vision instance, using the out-of-band management interface, in order to:
 - a. export traffic metadata in IPFIX format.
 - b. receive provisioning settings over HTTP.
2. It also requires bidirectional IP communication over the in-band interfaces with the Arista switches that compose an MSS Security Domain, in order to:
 - a. receive monitored traffic from the switches over L2oGRE.
 - b. maintain the healthiness of L2oGRE tunnels.


In-band Connectivity Requirements

The ZTX Monitor Node throughput is used for extracting and exporting communication session metadata from mirrored traffic received by the switches that compose one - and one only - Security Domain. The traffic received by the ZTX consists of a copy (mirror) of the original layer-2 traffic packets that match one or more monitor rules, truncated to the first 256 or 192 bytes¹ and embedded into a L2GRE envelope, which adds 24 bytes of overhead.

¹ This value is TOR hardware dependent and is not user configurable.

Centralized Configuration

Figure 3-12: ZTX Topology Centralized Design



- **Centralized:** The ZTX is placed within a service pod and its interfaces are physically connected to a pair of switches that have IP connectivity to the Security Domain.

This option is based on the best practice of using a dedicated self-contained unit in the network to provide all required network services with optimal scalability, efficiency and simplicity.

With this approach, the planned network throughput of the service pod, needs to account for the theoretical inbound capacity of the ZTX, unless further rate limiters are applied in the data path from the Security Domain.

The physical links of the ZTX that connect to the same physical device or, in case the peer devices use multi-chassis link aggregation, to the same device pair, can be grouped in a port channel. This choice may depend also on the selected routing design, which is discussed next.

IP Routing Requirements

To satisfy its in-band communication requirements, the ZTX must be provisioned with a distinct IP address, which needs to be reachable by the switches part of the Security Domain in order to create L2GRE tunnels that use the ZTX address as destination. The IP address of the ZTX can be advertised to the adjacent switches via dynamic routing, or the adjacent switches can be configured with a specific static route that the adjacent switches can redistribute in their current routing protocol.

From an outbound routing direction perspective, the ZTX requires in its routing table either every loopback address of the switches of the Security Domain or an aggregate prefix that includes them. This information is needed for the sole purpose of a GRE tunnel health check, and can be also provided via dynamic or static routing.

In case the ZTX is physically adjacent to a pair of switches that support MLAG, the choice of using dynamic vs. static routing influences the decision on how to bundle the ZTX interfaces into one or two port channels.

VRF and Security Zone Awareness

Monitoring rules are elements of MSS policies, configured on the Security Domain switches, which apply to a specific VRF, representing a determined security zone. On the ZTX, the VRF is identified via a GRE keyword field, present in the mirrored packet, and does not require to be declared in the device configuration.

L2GRE Tunnel Provisioning

The provisioning of L2GRE tunnels on the ZTX and the Security Domain switches is automatic and triggered by creating or modifying monitoring rules and other MSS objects in CloudVision.


3.2.2.2 Managing ZTX with CloudVision

Once the ZTX device has been onboarded to CloudVision and the in-band communication with the Security Domain is complete, the next step is to associate it with a MSS Monitor Object in CloudVision, so it can be referenced by one or more policy rules.

An MSS Monitor Object is a structure that defines how a ZTX device can communicate with a Security Domain.


The definition of a Monitor Object is possible from the MSS Service, which, once enabled in General Settings, is available under **Network Services** in the Studios pane, as shown in the screenshot below.

Figure 3-13: CloudVision - MSS Service Studio Selection




As for any other studio, before making any edit to the MSS objects, a change-control workspace needs to be created, as shown in the following image.

Figure 3-14: Create Workspace



The studio form presents most of the MSS objects structured in a multi-level tabular format. Monitor Objects are listed in a dedicated table in the bottom portion of the studio. A new object can be defined using **Add Monitor Object**, as shown in the following image:

Figure 3-15: Add Monitor Object



Parameter	Description	Recommended Value
Name	Unique string identifying the Monitor Object that can be referenced by a policy rule	
Monitor Node	Associated ZTX device among those present in the device inventory	
Exporter Interface	Interface name on the ZTX that is used as IPFIX exporter and L2GRE tunnel termination	
Active Timeout	Active timeout period in ms for exporting IPFIX reports	long-term setting: 30000 – 300000 ms temporarily for initial deployment: 3000 - 30000 ms
Tunnel Destination IP	IP address of the ZTX used as L2GRE tunnel destination on switches part of the Security Domain	
Tunnel Source Interface	The interface name on the Security Domain switches is used as an L2GRE tunnel source.  Note: All switches consistently use the same interface name.  Important: Tunnel Source Interface IP should be unique to each TOR. In the MLAG deployment scenario, specify the Tunnel Source Interface IP unique to each MLAG peer device rather than specifying interfaces with shared IP addresses.	
Truncation	Boolean field, indicating if mirrored traffic is truncated or not	Yes
Rate Limit	Rate limiter expressed in Mbps applied on Security Domain switches to mirrored traffic per VRF sent to the ZTX	10,000

Once a new Monitor Object entry has been populated with all required values, it is possible to select it inside the field **Monitor Name** in multiple policy rules part of the same policy.

The same Monitor Object can be concurrently referred to by multiple policies, while a policy (associated to a security zone) can only use one Monitor Object.

First, it is necessary to navigate in the studio form to the **Policies** table, and from there, by clicking on the desired policy entry in the **Rules** column, it is possible to view and edit the corresponding policy rules. The following two images are provided as a reference.

Figure 3-16: CloudVision - Policies Table

The screenshot shows the 'Policies' table in the CloudVision MSS Service interface. The table has the following data:

Name	Description	Rules
Policy_Traffic_Rule1	Policy_Traffic_Rule1	Rule1
Policy_Traffic_Rule2	Policy_Traffic_Rule2	Rule2

Figure 3-17: CloudVision - Rules Table

The screenshot shows the 'Rules' table in the CloudVision MSS Service interface. The table has the following data:

Name	Description	Source	Destination	Services	Action Type
Rule1	Other rule	10.0.0.0/24	10.0.0.0/24	HTTP, HTTPS	Allow All
Rule2	Other rule	10.0.0.0/24	10.0.0.0/24	HTTP, HTTPS	Allow All
Rule3	Other rule	10.0.0.0/24	10.0.0.0/24	HTTP	Allow All
Rule4	Other rule	10.0.0.0/24	10.0.0.0/24	HTTPS	Allow All
Rule5	Other rule	10.0.0.0/24	10.0.0.0/24	HTTPS	Allow All
Rule6	Other rule	10.0.0.0/24	10.0.0.0/24	HTTP	Allow All
Rule7	Other rule	10.0.0.0/24	10.0.0.0/24	HTTPS	Allow All
Rule8	Other rule	10.0.0.0/24	10.0.0.0/24	HTTP	Allow All
Rule9	Other rule	10.0.0.0/24	10.0.0.0/24	HTTPS	Allow All
Rule10	Other rule	10.0.0.0/24	10.0.0.0/24	HTTP	Allow All

Once a new Monitor Object entry has been created, it is necessary to validate its parameters by clicking on **Autofill** on the studio form, as shown in the following image:

Figure 3-18: CloudVision - Autofill



If no errors are reported, the change-control workspace can be reviewed, using **Review Workspace**, and subsequently executed, following the same workflow used by other studios.

Operational Aspects

This section is designed to convey the basic knowledge necessary to effectively interact with the MSS system. It focuses on practical MSS features and best practices for ongoing system management.

MSS Dashboard

The [MSS Dashboard](#) provides you with an at-a-glance overview of your network security domain or domains. View statuses of data sources, MSS devices, and monitoring nodes, as well as recent events, network topology, and more.

Policy Manager

Use [Policy Manager](#) to view the elements of security domain policies, including domains, policies, groups, and services. Review and accept dynamic groups that CloudVision has learned from onboarded data sources.

Policy Monitor

Use the [Policy Monitor](#) to view policy statistics, including packet and byte counts for security policies that have been pushed to devices.

Policy Builder

Use the [Policy Builder](#) to generate policy rule recommendations that can be reviewed, edited, deleted, submitted, or archived. Once you submit new policy rules, configuration changes are pushed to all devices in the relevant security domain VRF.

Policy Logs

[Policy Logs](#) provides an at-a-glance view of recent changes to externally-configured MSS groups. Logs are generated when a group from an onboarded data source is added, deleted, or changed at the source. Changes include member prefixes being added or removed.

MSS Studio

Use the [MSS Studio](#) to define static groups and policy rules for your security domain or domains and to configure traffic monitoring that allows CloudVision to generate security policy recommendations in the Policy Builder.


Create a Zero Trust Network with MSS

CloudVision's Multi-Domain Segmentation Service (MSS) gives you the building blocks to easily implement a [Zero Trust Network](#). Zero trust networking minimizes lateral movement by segmenting the network into increasingly smaller perimeters. Within these, communication between endpoints is governed by custom forwarding rules and only approved connections are permitted.

4.1 MSS Dashboard

The MSS dashboard provides you with an at-a-glance overview of your network security domain or domains. View statuses of data sources, MSS devices, and monitoring nodes, as well as recent events and network topology. Recent Logs displays changes to static and dynamically-learned groups. Use the Policy Monitor table to quickly view the impact of security policy rules on your network in packet and byte counts.

Figure 4-1: MSS Dashboard



Group Sources

The Group Sources panel displays the streaming status of data sources that CloudVision uses to dynamically discover groups.

Click **View in Devices** to see data source details under [Device Registration](#).

Figure 4-2: Group Sources



Monitoring Nodes

Monitoring Nodes shows you the streaming status of monitoring nodes that have been onboarded to CloudVision.

Figure 4-3: Monitoring Nodes



Click **View in Policy Builder** to view the status of monitoring rules that are using the nodes.


MSS Devices

View device capacity for devices that have been tagged with the `security-domain: tag` label.



Note: Device capacity does not represent the average of TCAM utilization across all device chips. It represents capacity for the chip with the highest utilization.

Figure 4-4: MSS Devices



Click **View in MSS Domains** to view details on device LPML and TCAM utilization rates in [Policy Manager](#). Utilization is provided both as a percentage and as a count of total entities used.

Events

The events panel displays events that are associated with devices that have been tagged with the **security-domain:** tag label or devices that are streaming traffic-policy configurations. Use the severity icons at the top of the panel to filter events by severity.

Figure 4-5: Events




Use the severity icons at the top of the panel to filter events by severity. Click **View in Events** for further event details.

Topology View

View a topology of devices that have been tagged with the security-domain: tag label. The **Topology View** panel displays a topology including numerous devices. Use the icons to zoom to fit, expand, or collapse the topology view.

Figure 4-6: Topology View



Policy Logs

Policy Logs highlights changes to static and dynamically-learned groups. The panel displays the 50 most recent changes from those that were made in the last 7 days.

Policy Monitor

The **Policy Monitor** table provides packet and byte counts for security policy rules that have been pushed to devices.



Note: Counters shown in the Policy Monitor table are for traffic across all VRFs where the rule has been applied.

The **Policy Monitor** table shows 5 policy rules with relevant details, including packet and byte counts. Selecting a device name will take you to the device's page in Inventory. Selecting a **Policy**, **Rule**, or **Security Domain** name will take you to the relevant tab in the Policy Manager.

Figure 4-7: Policy Monitor

Policy	Rule	Source	Destination	Security Domain	Device Name	Packets	Bytes
pol-sd4-v1	rulem	<any>	<any>	domain1	Id579	1.5G	1.5 TB
pol-sd4-v2	rulem2	<any>	<any>	domain1	Id579	10.1G	10.1 TB
pol-sd4-v4	rule18	+1 ...	+1 ...	domain1	Id579	0	0.0 B
pol-sd4-v4	rule1-reverse	+1 ...	+1 ...	domain1	Id579	0	0.0 B
pol-sd4-v4	rule32-reverse	+1 ...	+1 ...	domain1	Id579	0	0.0 B

4.2 Policy Manager

Policy Manager is where you'll view the elements of security domain policies, including domains, policies, and groups. You'll also review and accept dynamic groups that CloudVision has learned from onboarded data sources.

View Security Domain Policies

The Policy Manager is made up of five tables where you can view and sort domains, policies, rules, groups, and policy objects such as services, traffic monitors, and source and destination IP groups.

Domains

The Domains table includes a list of devices from your CloudVision device inventory. It provides details on device LPML and TCAM utilization rates. Utilization is provided both as a percentage and as a count of total entities used.

Note: The utilization percentage is not an average of capacity across all device chips. It represents capacity for the chip with the highest utilization. View utilization per device chip in [Device Hardware Capacity](#).

Figure 4-8: Policy Manager - Domains



Device	Model	IP Profile	LPM Utilization	LPM Entities	TCAM Utilization	TCAM Entities	Domain
egp583	7265CKM-325	Unassigned	2%	387	8%	793	egp-ent1
gln478	7265CR26-30	Unassigned	0%	88	0%	0	Unassigned
in442	278-73826-185	Unassigned	0%	0	0%	0	Unassigned
in488	278-73826-185	Unassigned	0%	0	0%	0	Unassigned
l3d79	7000CK3-325	Unassigned	2%	387	8%	796	storage-1
l3d81	7000CK3-325	Unassigned	2%	387	8%	796	storage-1
l3d88	7265CKM-325	Unassigned	0%	14	0%	0	Unassigned
l3d87	7265CR36-325	Unassigned	0%	3	0%	0	Unassigned
l3d81	7265CR36-325	Unassigned	1%	8	0%	72	storage-1
l3d88	7265CR36-325	Unassigned	8%	1,808	0%	14	storage-1

Devices that have been assigned to a security domain in the [MSS Studio](#) will include the domain name in the **Domain** column, while unassigned devices will be designated as Unassigned. ZTX monitor appliances will show as **Unassigned**.


Use the drop-downs to sort the table by security domain and by device. See relevant security domain details by clicking on the row of an assigned device.

Policies

The Policies table includes a list of security policies configured in the [MSS Studio](#). It provides an overview of the number of devices and rules that are included in each policy.

Policies designated as Assigned have been assigned to a security domain, while devices designated as Unassigned have not.

Figure 4-9: Policy Manager - Policies



The screenshot shows the 'Policy Manager' interface with a navigation bar at the top containing 'Domains', 'Policies', 'Rules', 'Groups', and 'Policy Objects'. Below the navigation bar are three filter dropdowns: 'Domain: Any', 'VRF: Any', and 'Assignment: Any'. The main table displays a list of policies with the following columns: Name, Rules, Devices, Domain, VRF, and Assigned. The 'Assigned' column uses green labels for 'Assigned' and red labels for 'Unassigned'.

Name	Rules	Devices	Domain	VRF	Assigned
pol-e22	1	1	egw-rn1	default	Assigned
pol-e21	1	0			Unassigned
pol-e23	1	0			Unassigned
pol-e24-v1	1	1	domain1	VRF-1	Assigned
pol-e24-v2	1	1	domain1	VRF-2	Assigned
pol-e24-v3	1	1	domain1	VRF-3	Assigned
pol-e24-v4	34	1	domain1	default	Assigned
pol-e25	1	0			Unassigned

Use the drop-downs to sort the table by security domain, VRF, and whether or not policies have been assigned.

View a list of the rules included in a policy by clicking on the row of an assigned policy.

Figure 4-10: Policy Manager - Rule

The screenshot shows the 'Policy Manager' interface. The 'Policies' tab is selected. A table lists various policies with columns: Name, Rules, Devices, Domains, VRF, and Assigned. The 'pol-sd4-v4' policy is highlighted. To the right, a detailed view of 'pol-sd4-v4' is shown, including a list of rules (rule1, rule2, rule3) and their configurations (Source, Destination, Service, Action, Direction). The 'rule1' configuration is expanded, showing 'Source' as 'any', 'Destination' as 'any', 'Service' as 'all', 'Action' as 'forward', and 'Direction' as 'Bi-directional'.

Name	Rules	Devices	Domains	VRF	Assigned
pol-sd1	1	1	egp-only	default	Assigned
pol-sd1	1	0			Unassigned
pol-sd3	1	0			Unassigned
pol-sd4-v1	1	1	domain1	VRF-1	Assigned
pol-sd4-v2	1	1	domain1	VRF-2	Assigned
pol-sd4-v3	1	1	domain1	VRF-3	Assigned
pol-sd4-v4	24	1	domain1	default	Assigned
pol-sd5	1	0			Unassigned

Rules

The Rules table includes all security policy and monitoring rules, including those that were statically configured in the [MSS Studio](#) and those that were recommended by CloudVision and accepted as part of a security policy.

Figure 4-11: Policy Manager - Rules

The screenshot shows the 'Policy Manager' interface with the 'Rules' tab selected. A table lists rules with columns: Rule, Policy, Source, Destination, Service, Action, and Direction. The 'rule1' row is highlighted, showing its configuration details.

Rule	Policy	Source	Destination	Service	Action	Direction
rule1	pol-sd1	any	any	any	drop-and-monitor	→
rule2	pol-sd1	any	any	any	forward-and-monitor	→
rule3	pol-sd3	any	any	any	forward-and-monitor	→
rule4	pol-sd4-v1	any	any	any	forward-and-monitor	→
rule5	pol-sd4-v2	any	any	any	forward-and-monitor	→
rule6	pol-sd4-v3	any	any	any	forward-and-monitor	→
rule7	pol-sd4-v4	any	any	any	forward	↔

Use the drop-downs to sort the table by security policy and forwarding action.

Select any row in the table to view detailed lists of the policy objects included in the rule and the security policies that the rule is used in.

Groups

The Groups table includes a list of all groups, including those that were statically configured in the MSS Studio and those that were dynamically learned by external data sources and accepted for use in security policy configuration. It provides an overview of the group source, category, and number of members.

Groups designated as Active are being used in security policies, while groups designated as Accepted are not.

Figure 4-12: Policy Manager - Groups

Policy Manager

Domains

Policies

Rules

Groups

Policy Objects

IP Address

Search IP Address

Filters

Source: Any

Category: Any

Manage Static Categories

Review Groups

Name ↓	Source	Categories	Members	Status
Enw62	CSV	Site	19	Active
Enw61	CSV	Site	19	Active
Enw60	CSV	Site	19	Active
Enw6	CSV	Site	19	Active
Enw66	CSV	Site	19	Active
Enw666	CSV	Site	2	Accepted
Enw667	CSV	Site	19	Accepted
Enw668	CSV	Site	19	Accepted

Use the drop-downs to sort the table by group source and category.

View a list of group members and any policy rules that the group is included in by clicking on the row of an active policy.

Figure 4-13: Policy Manager - Group Members

Policy Manager

Domains

Policies

Rules

Groups

Policy Objects

IP Address

Search IP Address

Filters

Source: Any

Category: Any

Manage Static Categories

Review Groups

Name ↓	Source	Categories	Members	Status
Enw62	CSV	Site	19	Active
Enw61	CSV	Site	19	Active
Enw60	CSV	Site	19	Active
Enw6	CSV	Site	19	Active
Enw66	CSV	Site	19	Active
Enw666	CSV	Site	2	Accepted
Enw667	CSV	Site	19	Accepted
Enw668	CSV	Site	19	Accepted
Enw669	CSV	Site	19	Accepted
Enw669	CSV	Site	19	Accepted
Enw669	CSV	Site	19	Accepted

Enw62

Members 19 (19)

192.168.1.10/32

192.168.1.10/32

192.168.1.10/32

192.168.1.10/32

192.168.1.10/32

192.168.1.10/32

192.168.1.10/32

192.168.1.10/32

Rules 2 (2)

> permit_log_CSV-Enw62_CSV-App2

> permit_log_CSV-Enw62_CSV-App2

Policy Objects

The Policy Objects table includes a list of policy building blocks including destination IP groups, source IP groups, services, and traffic monitors. It provides an overview of the type of policy objects as well as their details and the number of policies that they're included in.

Figure 4-14: Policy Manager - Policy Objects



Object Name	Type	Details	Policy Count
CDVT-Den	Destination IP Group	10.10.10.10/24 - 10.10.10.10/24	1
CDVT-Prod	Source IP Group	10.10.10.10/24 - 10.10.10.10/24	1
s1	Service	Dest Port 80, Source Port all, Protocol tcp	1
s2	Service	Dest Port 8, Source Port all, Protocol udp	1
s3	Service	Dest Port 443, Source Port all, Protocol tcp	1
s4	Service	Dest Port 3784, Source Port all, Protocol udp	1
s5	Service	Dest Port 80, Source Port all, Protocol udp	1
mon1	Traffic Monitor	Active	1
mon2	Traffic Monitor	Active	1
mon3	Traffic Monitor	Active	1
mon4	Traffic Monitor	Active	1
mon5	Traffic Monitor	Active	1

Use the drop-downs to sort the table by security policy and policy object type. Click on any row in the table to view a detailed list of the policies that the object is included in.

Figure 4-15: Policy Manager - Policy Objects Details



Object Name	Type	Details	Policy Count
CDVT-Den	Destination IP Group	10.10.10.10/24 - 10.10.10.10/24	1
CDVT-Prod	Source IP Group	10.10.10.10/24 - 10.10.10.10/24	1
s1	Service	Dest Port 80, Source Port all, Protocol tcp	1
s2	Service	Dest Port 8, Source Port all, Protocol udp	1
s3	Service	Dest Port 443, Source Port all, Protocol tcp	1
s4	Service	Dest Port 3784, Source Port all, Protocol udp	1
s5	Service	Dest Port 80, Source Port all, Protocol udp	1
mon1	Traffic Monitor	Active	1
mon2	Traffic Monitor	Active	1
mon3	Traffic Monitor	Active	1
mon4	Traffic Monitor	Active	1
mon5	Traffic Monitor	Active	1

mon4

Details

- Node:** W7824220011
- Egress Interface:** 10.10.10.10
- Active Timeout:** 30s
- Tunnel Destination IP:** 10.10.10.10
- Tunnel Source Interface:** 10.10.10.10
- Rate Limit:** 100,000 mbps
- Mode:** Top

Review and Accept Dynamic Groups

Once you've [onboarded data sources](#) to CloudVision, you'll be able to review dynamically-learned groups in the Policy Manager.

1. From the Groups tab, click **Review Groups**.

Figure 4-16: Review Groups



This launches a modal where you can review dynamically-learned groups.

2. Enable the checkbox next to groups to accept them.


Figure 4-17: Review Dynamic Groups



Tip: Clicking on any row in the table will allow you to view group members.

Groups that are not selected initially are categorized as Ignored. Ignored groups can be accepted during a subsequent review provided that the review table is filtered to show All Statuses or Ignored, as shown below:

Figure 4-18: Review Dynamic Groups - Ignored



3. Click **Submit Changes**.
4. Once the changes have been validated, click **Finish** or **Configure Additional Inputs** to accept additional groups.

Figure 4-19: Review Dynamic Groups



If validation fails, you'll have the option to manually review the changes.

Accepted groups will appear in the Groups table in Policy Manager and will be available for configuring rules in the [MSS Studio](#).




Note: You can perform the same workflow of reviewing and accepting groups using the **Review Dynamic Groups** quick action in the MSS Studio.

4.3 Policy Monitor

You'll use Policy Monitor to view policy statistics, including packet and byte counts for security policies that have been pushed to devices.

Figure 4-20: Policy Monitor



The screenshot shows the 'Policy Monitor' interface with a table of policy statistics. The table has columns for Policy, Rule, Source, Destination, Security Domain, Device Name, Packets, and Bytes. The first two rows show data for 'pol-udA-v0' and 'pol-udA-v1' policies, while the last three rows show zero counts for 'pol-udA-v2', 'pol-udA-v3', and 'pol-udA-v4'.

Policy	Rule	Source	Destination	Security Domain	Device Name	Packets	Bytes
pol-udA-v0	rule0	10.10.10.1	10.10.10.2	Domain1	R101	10,000	10,000 B
pol-udA-v1	rule1	10.10.10.1	10.10.10.2	Domain1	R101	1,000	1,000 B
pol-udA-v2	rule2	10.10.10.1	10.10.10.2	Domain1	R101	0	0 B
pol-udA-v3	rule3	10.10.10.1	10.10.10.2	Domain1	R101	0	0 B
pol-udA-v4	rule4	10.10.10.1	10.10.10.2	Domain1	R101	0	0 B

Counters shown in the Policy Monitor are for traffic across all VRFs where the traffic policy has been applied.

Filter the table by device name, security domain, source, destination, rule, or by packet or byte count.

Clicking on a policy, rule, or security domain name will take you to the corresponding table in the [Policy Manager](#). Clicking a device name brings you to the device's page in [Inventory](#).

4.4 Policy Builder

You'll use the Policy Builder to generate and review security policy recommendations. After a monitoring rule is created for a security policy in the MSS Studio, the monitor node will conduct stateful traffic analysis. It sends session data back to CloudVision in order to map sessions to group-to-group communication. You'll then be able to generate policy rules that can be reviewed, edited, deleted, submitted, or archived to be revisited at a later time. Once you submit new policy rules, configuration changes will be pushed to all devices in the relevant security domain VRF.

View Monitoring Rules

All monitoring rules that you have configured in the [MSS Studio](#) are available to view in Policy Builder.

Figure 4-21: View Monitoring Rules

Rule ID	Name	VRF	Source	Destination	ZSS Node
60110670	Collecting sessions	domain1	VRF-0	<range> 27 </range>	WTW04220007 100,000 Mbps
43605796	Collecting sessions	domain1	default	<range> 27 </range>	WTW04220007 100,000 Mbps
60120769	Collecting sessions	edge-vrf0	default	<range> 27 </range>	WTW04220006 100,000 Mbps
436	Collecting sessions	domain1	VRF-0	<range> 27 </range>	WTW04220007 100,000 Mbps
61762962	Completed	edge-vrf0	VRF-1	<range> 27 </range>	WTW04220006 100,000 Mbps
60276230	Completed	domain1	VRF-1	<range> 27 </range>	WTW04220007 100,000 Mbps

Rules are categorized in the Monitoring Rules table according to their progress as follows:

- **Collecting Sessions:** CloudVision is collecting sessions via the monitoring node
- **Generating Rules:** In this temporary state, CloudVision is mapping sessions to group-to-group communication to generate security policy rule recommendations
- **Review Rules:** Rule recommendations are ready for review and can be accepted, accepted with edits, or rejected
- **Completed:** Policy recommendations have been reviewed and decisions made regarding implementation of the recommended rules



Tip: Monitoring rules marked as completed are hidden from the table by default but can be viewed by enabling the Show Completed toggle shown above.

Click on any monitoring rule in the Policy Builder table, to view monitoring rule details and progress.

Figure 4-22: Monitoring Rule



Optionally, click the Edit to rename the monitoring rule.

Monitoring Rule Details

Click **Details** under the **Policy** tab.

Figure 4-23: Monitoring Rule Details



This opens a drawer that includes the monitoring rule and monitoring session details.

Figure 4-24: Monitoring Rule Details



Monitoring Rule Progress

Monitoring rule progress is shown in the **Policy** tab, including the duration that the monitoring node has been collecting traffic sessions. In the below image, for example, the node has been collecting sessions for 13 hours and 46 minutes.

Click the Collected Session tab to view the sessions.

Figure 4-25: Collected Sessions



A summary of all collected sessions highlights the total number of sessions, as well as the number of sessions that are complete and half open. A half-open session indicates that only one-way communication has been observed and there has been no response yet from the destination IP address.

Individual sessions are listed in the table with forward and reverse pack and byte counts.

Figure 4-26: Individual Sessions

Policy Builder

ccf10b70

Monitoring Rule

Policy: [Collected Sessions](#)

Total: 0

Complete: 0

Half Open: 0

Protocol: Any ▾ Source IP Address: Any ▾ Dest. IP Address: Any ▾ Dest. Port: Any ▾ Filter Columns

Source IP Address	Destination IP Address	Protocol	Destination Port	Forward Packets	Reverse Packets	Forward Bytes	Reverse Bytes	Complete	Half Open
192.168.1.10	192.168.1.10	UDP	700	1,336,771	2,241,423	1,473,822,239	3,182,144,646	0	0
192.168.1.10	192.168.1.10	UDP	700	1,336,528	2,271,534	1,473,188,762	3,227,576,878	0	0
192.168.1.10	192.168.1.10	UDP	700	1,337,223	2,224,488	1,474,384,384	3,179,677,000	0	0
192.168.1.10	192.168.1.10	UDP	700	1,006,719	2,288,086	1,473,828,862	3,237,746,098	0	0
192.168.1.10	192.168.1.10	UDP	700	1,336,714	2,258,294	1,473,878,482	3,206,863,840	0	0
192.168.1.10	192.168.1.10	UDP	700	1,336,867	2,257,478	1,473,812,748	3,207,778,354	0	0
192.168.1.10	192.168.1.10	UDP	700	1,336,330	2,256,475	1,473,182,888	3,204,788,518	0	0
192.168.1.10	192.168.1.10	UDP	700	1,336,333	2,256,313	1,473,183,874	3,206,687,808	0	0
192.168.1.10	192.168.1.10	UDP	700	1,006,719	2,242,586	1,473,825,868	3,181,736,638	0	0
192.168.1.10	192.168.1.10	UDP	700	1,337,307	2,282,574	1,473,848,878	3,184,253,988	0	0

Showing 10 rows

Generate Security Policy Rules

Once a monitoring rule has collected sessions, you'll be able to generate security policy rule recommendations.

1. Click on a monitoring rule that is collecting sessions.
2. Click **Generate Rules**.

Figure 4-27: Generate Rules




Note: The monitor node will continue to collect sessions for a configured monitoring rule until the rule has been deleted from the domain policy in the [MSS Studio](#). The **Generate Rules** button is enabled once a single session has been collected. You can view the elapsed time for session collection under the Policy tab, shown above, and the sessions themselves under the **Collected Sessions** tab in order to determine when to generate rules.

3. Select categories to generate security policy rules by enabling the relevant checkboxes in the modal.

Categories and group names are learned from [onboarded data sources](#).

Figure 4-28: Groups



Tip: Clicking on a row in the **Generate Rules** modal will enable you to view the address groups included in the category, as shown above.

4. Click **Generate**.




Note: On average it takes 5-10 minutes for CloudVision to generate security policy rules, but could take longer.

Review Security Policy Rules

Once CloudVision has generated security policy recommendations for a monitoring rule, it will be reclassified in the Monitoring Rules table and “Review rules” will appear below the Rule ID.

You can click on any monitoring rule with this designation to review security rule recommendations.

Figure 4-29: Review Rules



1. Review security rule recommendations.

You'll use the available drop-downs and icons to remove or edit rule recommendations. Drag and drop rules to reorder them in relation to each other and to existing security policy rules.

As shown below, some rules may require you to select an appropriate service, action, direction, or a combination of these in order to submit it.

Figure 4-30: Review Rule - Edit



Icons that appear next to the rule order number indicate whether the rule was recommended, recommended and edited, or created in the [MSS Studio](#).

Figure 4-31: Review Rule - Edit



Note: While editing rules you can also reset changes or regenerate rules. You'll regenerate rules to change the categories you're generating security rules for.

2. Click **Submit Rules**.

Figure 4-32: Submit Rules



Note: After submitting the rules, Policy Builder archives the session. View archived sessions using **Show Completed Sessions**.


This creates a change control that is auto-approved and auto-executed, pushing the new configuration to devices in the security domain VRF.

3. Optionally, click **Archive** to mark the rule as Completed and store it. Archived rules can be viewed but are not actionable.

4.5 Policy Logs

Policy Logs provides an at-a-glance view of recent changes to externally-configured MSS groups. Logs are generated when a group from an onboarded data source is added, deleted, or changed at the source. Changes include member prefixes being added or removed.

Figure 4-33: Policy Logs




The screenshot displays the 'Policy Logs' interface. At the top, there are filter buttons for 'Timestamp', 'Source Type', 'Source Name', 'Severity', and 'Message'. Below the filters, a table lists recent events. The table has five columns: 'Timestamp', 'Source Type', 'Source Name', 'Severity', and 'Message'. The data rows show events from April 8, 2025, including group changes, deletions, and additions. A timeline at the bottom of the interface shows the progression of time from 12:00 to 9:00 on April 8, 2025.

Timestamp	Source Type	Source Name	Severity	Message
Apr 8, 2025 11:30:07	External Source	MSSCEVDS	Info	1 group changed
Apr 8, 2025 11:30:07	External Source	MSSCEVDS	Info	2 groups deleted
Apr 8, 2025 09:28:44	External Source	MSSCEVDS	Info	2 groups added
Apr 7, 2025 12:58:08	External Source	MSSCEVDS	Info	1 group changed
Apr 7, 2025 12:13:55	External Source	MSSCEVDS	Info	12 groups added
Apr 7, 2025 12:13:55	External Source	MSSCEV1	Info	12 groups added

Filter the table by **Timestamp**, **Source Type**, **Source Name**, **Message**, or **Severity**.

Selecting a table row opens a side panel with log details, including the relevant group or groups and the associated changes. View logs for a selected time period by using the **Timepicker**.

Figure 4-34: Log Details



Timepicker

Use the timepicker to adjust the time frame for viewing historical data. Use the predefined time interval or customize it to show broader or more granular information.

The timepicker can be found at the bottom of pages that display metrics. Use it to view device information, dashboard metrics, or broader topology data.

Adjusting the Time

By default the timepicker shows you information in live time. Selecting **Live** returns the pointer to the current time.

4.6 MSS Studio

Use the [MSS Studio](#) to define static groups and policy rules for your security domain or domains and to configure traffic monitoring that allows CloudVision to generate security policy recommendations in the Policy Builder.

4.7 Create a Zero Trust Network with MSS

CloudVision's Multi-Domain Segmentation Service (MSS) gives you the building blocks to easily implement a zero trust network. Zero trust networking minimizes lateral movement by segmenting the network into increasingly smaller perimeters. Within these, communication between endpoints is governed by custom forwarding rules and only approved connections are permitted.

Leverage MSS tools like the Policy Manager, Policy Builder, and MSS Studio to make the zero trust microsegmentation and enforcement process simple and effective. CloudVision MSS automates the management of microperimeters by connecting to external sources and dynamically identifying and tagging endpoints and workloads. The service maps all communications within the network, giving you complete visibility into existing traffic flows. The observed traffic map becomes the basis for security policy recommendations that permit only trusted communications. MSS pushes zero trust policies to EOS devices, which distribute enforcement themselves or redirect traffic to a third-party firewall. Once the zero trust policies are deployed, MSS can monitor for policy violations and new traffic to help you keep your network safe and traffic rules up-to-date.

In the workflow outlined here, you'll use the Policy Builder to monitor traffic and incrementally build an allow list. Start by defining coarse security policy rules and move toward implementing more granular ones. Once the allow list is established, you'll implement zero trust by creating a rule to monitor and drop all east-west traffic that is not explicitly permitted by the security policy rules.

For more on Arista's zero trust solution, see [Arista Zero Trust Security for Cloud Networking](#).

4.7.1 Prerequisites

The Policy Builder generates policy rule recommendations to help you build the allowlist for your zero trust network. Before you can review recommendations, you'll need to complete some preliminary steps using [MSS Studio](#) and [Policy Manager](#).



Tip: See [Getting Started with MSS](#) for additional details on completing prerequisites.

1. Onboard the [ZTX Monitor Node Appliance](#) and register it in the [Inventory and Topology Studio](#).



Note: If the End-to-End Provisioning toggle is enabled in **General Settings > Features**, newly-onboarded devices will automatically be available to be registered for use in Studios.

2. Onboard any data sources that you'll use to define groups.
Supported data sources include Arista's network identity service AGNI, VMware vCenter, configuration management databases (CMDBs) like ServiceNow, and CSV files.
3. Accept groups discovered from onboarded data sources in the [Policy Manager](#).
4. Create a new or open an unsubmitted Studios [workspace](#) to configure a security domain, security policy, monitor object, and static groups in the [MSS Studio](#).
 - **Security Domain:** Enter a relevant security domain name in Security Domains and assign devices to the domain.
 - **Security Policy:** Create a policy name and associate it with the relevant VRF in Security Domains. You'll configure policy rules later.



Note: If you have multiple VRFs in your domain, you'll need to create a security policy for each VRF.

- **Monitor Object:** Enter relevant data about the onboarded ZTX-7250S appliance and connected TOR devices in order to open monitoring sessions between the monitoring node and EOS devices.



Note: MSS allows one monitor object per security domain. Therefore, multiple security domains require multiple monitor objects.

- **Static Groups:** Configure groups by IP prefix. To start with you must configure an "internal-networks" group that includes all internal subnets in the VRF.

You're now ready to create the monitoring rule that will allow the [Policy Builder](#) to generate policy rule recommendations.

4.7.2 Configuring Monitoring Rules

Create a rule that forwards mirrored east-west traffic to the monitor node. The monitor node will then begin collecting session data to create a traffic map. The Policy Builder will use the map to recommend policy rules relevant to your network.



Note: If there are multiple VRFs in your security domain, you'll configure a monitoring rule for each VRF.

1. Click **View** to configure rules for the policy that you associated with the domain VRF.

Figure 4-35: Security Domains



2. Click **Add Rule** to configure the monitoring rule for east-west traffic.



Note: As you continue to add rules, the east-west monitoring rule needs to remain the last rule in the policy rule list.

Figure 4-36: MSS Service ZTN Policy



- **Name:** Create a name to identify the monitoring rule.
 - **Description:** Optionally, enter a rule description.
 - **Source:** Select the `internal-networks` group that you configured in Static Groups.
 - **Destination:** Select the `internal-networks` group that you configured in Static Groups.
 - **Service:** Select `any`.
 - **Action:** Select `forward`.
 - **Direction:** Select `Unidirectional`.
 - **Monitor Name:** Select the name of the monitor object that you configured in Monitor Objects.
3. **Review and submit** the workspace to push configuration to all devices in the relevant security domain.

The ZTX monitor node will begin collecting session data from mirrored east-west traffic for you to review in the Policy Builder. North-south traffic not governed by the monitoring rule will be forwarded as normal.

4.7.3 Configuring Security Policy Rules


View sessions from traffic mirrored to the ZTX monitor node and begin to configure security policy rules for your zero trust network. You'll begin by generating, reviewing, and editing coarse rules and proceed incrementally to more granular ones. This typically means starting with rules to govern forwarding among environments and moving toward rules that manage traffic among applications.



Tip: View [Policy Builder](#) documentation for details on how to view collected sessions and generate and review policy rule recommendations.

As an example, imagine you have a security domain whose internal networks are made up of two environments, production and development, and three applications. These might be email clients, web browsers, or other software applications.

Figure 4-37: Security Domain Example



In this case, you'll have groups that represent the internal networks, the production environment, the development environment, and each of the three applications. More granularly, you'll have groups that represent endpoints that may be part of the same application, but different environments. Examples include, App1-Prod, App1-Dev, etc.

These groups will have been statically-defined in the [MSS Studio](#) or dynamically discovered using onboarded data sources and accepted in the [Policy Manager](#) for use in configuring security policy rules.

4.7.4 Generating Coarse Rules

Once the monitoring rule has collected sessions, you can generate rules. Begin by generating rules that regulate traffic between environments.


1. Click on the monitoring rule in the Policy Builder.
2. Click **Generate Rules**.

Figure 4-38: Generate Rules



3. Select the category that represents internal network environments.

Figure 4-39: Generate Rules Internal Network



4. Click **Generate**.

The Policy Builder will generate policy rule recommendations.

5. Use the icons, drop-downs, and input fields provided to edit and accept or delete policy rule recommendations.

Figure 4-40: Review & Edit Rules

	Source	Destination	Protocol	Dest Port	Service	Action	Direction
1	CSG1-Prod	CSG1-Prod	udp	334	Enter name	forward	Bi-directional
2	10.10.0.0/16	CSG1-Prod	udp	332	Enter name	forward	Bi-directional
3	10.10.0.0/16	CSG1-Prod	udp	330	Enter name	forward	Bi-directional

At this stage, edit and accept only broad environment rules.

Using the sample security domain, assume that you want to allow all endpoints in the development environment (11.11.0.0/16) to talk to one another, but you want to restrict traffic between development and production (10.10.0.0/16).

You can expect the [Policy Builder](#) to generate rules for production-production traffic, development-production traffic, and development-development traffic.

Tip: If the Policy Builder doesn't generate appropriate rule recommendations, click **Regenerate Rules** in the **Review & Edit Rules** header for new recommendations.

Figure 4-41: Remove Rule

	Source	Destination	Service	Action	Direction	
1	Production	Production	any	forward	Bi-directional	Remove Rule
2	Development	Production	any	drop	Bi-directional	
3	Development	Development	any	forward	Bi-directional	
4	Internal	Internal	any	forward	Uni-directional	

- **Production to Production:** Delete the production to production rule recommendation.
Allow the default forwarding rule for east-west traffic to forward this traffic. Later, you'll add additional granularity to production environment forwarding by configuring rules at the application level.
- **Development to Production:** Edit and accept the development to production rule recommendation to restrict traffic between the two environments by selecting the following inputs.
 - **Service:** any
 - **Action:** drop
 - **Direction:** Bidirectional
- **Development to Development:** Edit and accept the development to development rule recommendation to allow all endpoints to communicate freely with one another by selecting the following inputs.
 - **Service:** any
 - **Action:** forward

- **Direction:** Bidirectional
6. Order rules appropriately.
New environment rules should precede the internal_networks monitoring rule, as shown below.

Figure 4-42: Environment Rules



		Source	Destination	Service	Action	Direction
1		Development	Production	any	drop	Bi-directional
2		Development	Development	any	forward	Bi-directional
3		Internal	Internal	any	forward	Uni-directional

At this stage, only production-production traffic is mirrored to the ZTX monitor node in order to generate rules for the production environment with application-level granularity.

As rules are refined, traffic governed by rules with **higher precedence** than the monitoring rule is offloaded from the ZTX monitor node.

4.7.5 Generating Granular Rules

Once you've configured coarse, environment-level rules, you'll begin generating and editing more granular policy rule recommendations to govern traffic between applications.

1. Follow steps 1- 4 of the workflow for generating coarse policy rule recommendations, but select the categories that correlates to internal network applications.

Figure 4-43: Granular Rules

Generate Rules

Policy Details

Policy: pol-sd1 Domain: egv-only VRF: default

Generate Using ① 1 Category Checked

Category	Source	Groups
☒ App	CSV1	2
☒ Env	CSV1	2
☐ Service	CSV1	100
☐ Hostname	CSV1	100
☐ StaticCategory	static	0

Showing 5 of 5 rows

Select a category row to view its members

2. Just like before, you'll use the icons, drop-downs, and input fields provided to edit and accept or delete policy rule recommendations.

Using the sample security domain, assume that you want to govern production environment traffic to explicitly:

- Restrict App 1 from communicating with App 3 using https.
- Allow App 1 to communicate with App 2 using https.
- Allow App 2 to communicate with App 3 using https.
- Restrict App 3 from initiating communication with App 1 or App 2.

To do so, you'll want to edit and accept policy rule recommendations as shown below:



Note: If you've configured a service in the MSS Studio by defining its protocol and port, the [Policy Builder](#) will identify the service accordingly in its recommended rules. If a protocol and port has no configured service name, then you can provide one directly in the recommended rule by entering a name in the Service field.

Figure 4-44: Regenerate Rules

Review & Edit Rules generated 8 minutes ago Reset Changes Regenerate Rules

		Source		Destination		Service	Action	Direction	
1		Prod	App1	↔	Prod	App3	https	drop	Bi-directional
2		Prod	App1	↔	Prod	App2	https	forward	Bi-directional
3		Prod	App2	→	Prod	App3	https	forward	Bi-directional
4		Prod	App3	→	Prod	App2	any	drop	Uni-directional
5		Prod	App3	→	Prod	App1	any	drop	Uni-directional
6		Development		↔	Production		any	drop	Bi-directional
7		Development		↔	Development		any	forward	Bi-directional
8		Internal		↔	Internal		any	forward	Uni-directional

- **Production-App1 to Production-App 3:**
 - **Service::** https
 - **Action::** drop
 - **Direction::** Bidirectional
 - **Production-App1 to Production-App 2:**
 - **Service::** https
 - **Action::** forward
 - **Direction::** Bidirectional
 - **Production-App2 to Production-App 3:**
 - **Service::** https
 - **Action::** forward
 - **Direction::** Bidirectional
 - **Production-App3 to Production-App 2:**
 - **Service::** any
 - **Action::** drop
 - **Direction::** Unidirectional
 - **Production-App3 to Production-App 1:**
 - **Service::** any
 - **Action::** drop
 - **Direction::** Unidirectional
3. Order rules appropriately. CloudVision processes rules sequentially, so it's important to consider rule order. The default east-west monitoring rule should remain in the last position.

-
4. Incrementally create more granular rules to govern forwarding of east-west network traffic. As rules are refined, traffic governed by rules with higher precedence than the monitoring rule is offloaded from the ZTX monitor node. Once all rules are in place, you can enable your zero trust network.

4.7.6 Implementing Zero Trust

Implementing zero trust is an optional step that can be taken once you are confident that you've explicitly allowed all trusted traffic. A zero trust network grants explicit permission only to known trusted traffic and mirrors only untrusted traffic to the ZTX monitor node, which should be assessed on an ongoing basis in the Policy Builder.

Edit the monitoring rule to implement zero trust.

1. Change the action in the internal-networks monitoring rule. Replace `forward` with `drop`.
This drops, but continues to monitor, all traffic that isn't governed by forwarding rules.
2. Optionally, eliminate drop rules, as the default behavior for the monitoring rule is to drop all traffic that isn't explicitly allowed.

In this case, the final allowlist for the sample network would appear as follows:

Figure 4-45: Drop Rules

Review & Edit Rules generated 8 minutes ago Reset Changes Regenerate Rules

		Source	Destination	Service	Action	Direction
1		Prod App1 ↔ Prod App2	https	forward	Bi-directional	
2		Prod App2 → Prod App3	https	forward	Bi-directional	
3		Development ↔ Development	any	forward	Bi-directional	
4		Internal ↔ Internal	any	drop	Uni-directional	



Note: Eliminating drop rules creates a simple allowlist for network traffic, but increases traffic mirrored to the monitor node.

Once the allowlist is in place, continue to monitor dropped traffic mirrored to the ZTX monitor node. As needed, generate new rule recommendations in the Policy Builder to explicitly permit new, trusted traffic.

Troubleshooting MSS

This section contains valuable information on troubleshooting common issues, and perform routine maintenance tasks.

- [Show Commands](#)
- [Tracing](#)
- [Considerations](#)

5.1 Show Commands

The following **show** commands help with troubleshooting a ZTX Monitor Node.

Check the interface status to ensure that the port-channel and member ports are connected:

```
ZTX# show interfaces status
Port  Name  Status      Vlan   Duplex  Speed  Type           Flags  Encapsulation
Et1/1      connected in Po1 full   10G    10GBASE-SR
Et1/2      connected in Po1 full   10G 1  0GBASE-CR
..
Po1 connected routed full 40G N/A
```

Check each GRE tunnel interface status to ensure that the status is UP (multiple GRE tunnels may terminate on the interface):

```
ZTX# show interfaces tunnel 0
Tunnel0 is up, line protocol is up (connected)
Hardware is Tunnel, address is 0000.0000.0000
Tunnel source 10.10.254.1, destination 10.10.254.2
Tunnel protocol/transport GRE/IP
Hardware forwarding enabled
Tunnel transport MTU 1476 bytes (default)
Tunnel underlay VRF "default"
Up 3 days, 53 minutes, 22 seconds
```

Check the flow tracking feature status is active on all GRE monitor tunnels from TORs:

```
ZTX# show flow tracking firewall distributed
Flow Tracking Status
Type: Distributed Firewall
Running: yes, enabled by the 'flow tracking firewall distributed' command
Tracker: flowtrkr
Active interval: 300000 ms
Inactive timeout: 15000 ms
Groups: IPv4
Exporter: exp
VRF: default
Local interface: Loopback0 (10.10.254.1)
Export format: IPFIX version 10, MTU 9152
DSCP: 0
Template interval: 3600000 ms
Collectors:
127.0.0.1 port 4739
Active Ingress Interfaces:
Tu0, Tu1
```

Check if mirrored flows are seen at the ZTX Node:

```
ZTX# show firewall distributed instance session-table
Legend
eph - Ephemeral port
Sessions: 5
VRF      Proto Source/Destination Fwd/Rev Src VTEP IP      Fwd/Rev Pkts  Fwd/Rev Bytes Complete
Half-Open Start Time Destination
-----
vrf2     UDP    1.1.1.1:50004      10.10.254.2      1          428      0 1 2024-10-28
11:13:09 1.1.1.4:1001 10.10.254.3 1 428
vrf1     UDP    1.1.1.1:eph        10.10.254.2      5          2140     5 0 2024-10-28
11:13:09 1.1.1.3:1001 10.10.254.3 5 2140
vrf1     UDP    1.1.1.1:eph        10.10.254.2      5          2140     5 0 2024-10-28
11:13:09 1.1.1.2:1001 10.10.254.3 5 2140
```

If mirrored flows are not seen at the ZTX Node, check for drops:

```
ZTX# show platform sfe counters | nz
Name                               Owner                               Counter Type Unit    Count
-----
Tunnel-Global-gre_decap_drop_pkts  Ip4TunDemux                       module   packets 400
Tunnel-Global-tun_decap_drop_pkts  Ip4TunDemux                       module   packets 260
IpInput_Tunnel0-Stateful_drop_counter IpInput_Tunnel0                   module   packets 47
```

If mirrored flows are not seen at CloudVision Portal, check if ZTX Node has exported the flows, and if there are no failures in IPFIX export:

```
switch# show agent sfe threads flow cache scan counters
Purged count: 501
IPFIX export count: 354
IPFIX failed export count: 0
```

The following is a brief explanation of the output:

- **Type:** Distributed Firewall type for ZTX nodes.
- **Running:** “yes” indicates that the flow tracking feature is successfully running.
- **Tracker:** Name of flow tracker configuration.
- **Active interval:** Interval after which IPFIX data packet is exported for active sessions. Active interval is set to 1800000ms (30mins) by default and can be modified in MSS Studio.
- **Inactive timeout:** The time after which sessions are considered inactive if no packets are received. Inactive timeout is not configurable and defaults to 15000ms.
- **Groups:** Currently, only IPv4 packets are supported.
- **Exporter:** Name of exporter configuration.
- **VRF:** VRF used for IPFIX export.
- **Local Interface:** Local interface used for IPFIX export.
- **Export format:** IPFIX version and IP MTU used for exported IPv4 packets.
- **DSCP:** Differentiated Service Code Point value used in exported IPv4 packet header.
- **Template Interval:** Time interval between successive IPFIX template export to collector.
- **Collectors:** List of IPFIX Collector IP and Port. 127.0.0.1 indicates local IPFIX collector running on a ZTX device. The local IPFIX collector will send the exported flows to CVP.
- **Action Ingress Interfaces:** Displays all the tunnel interfaces on which IPFIX flow tracking is running.

5.2 Tracing

Enabling tracing can seriously impact the switch's performance in some cases. Please use it cautiously and seek advice from an Arista representative before enabling it in any production environments.

```
trace Sfe setting IpfixWalker*/*
```

5.3 Considerations

When deploying Multi-domain Segmentation Services (MSS) with the ZTX-7250S-16S Monitor Node, several crucial **technical aspects** must be considered to ensure optimal performance and policy enforcement.

- [ZTX-7250S-16S Monitor Node](#)
- [Self-IP Support](#)
- [Layer 2 Devices](#)
- [Network Address Translation](#)
- [Access Control Lists and Policy-Based Routing](#)

5.3.1 ZTX-7250S-16S Monitor Node

Session Capacity and Traffic Throughput

The **ZTX-7250S-16S Monitor Node** can handle a maximum of **32 million session entries** concurrently. Session entries include a mix of aggregate, short-lived ephemeral, and persistent non-ephemeral sessions. Monitoring your network's session count is vital to avoid exceeding this limit, which could impact performance or lead to dropped connections. Additionally, the node supports up to **80Gbps of incoming monitor traffic** when all 16 Ethernet ports are actively connected to upstream service Top-of-Rack (TOR) switches. Designing your network to leverage all available ports will maximize monitoring throughput.

5.3.2 Self-IP Support

MSS Studio automatically applies a **Self-IP rule** to permit unicast traffic destined for the device. While this is usually sufficient for control plane protocols, those relying on **multicast PDUs**—like PIM, OSPF, or IPv6 Neighbor Discovery for BGPv6—might fail to establish if a default **deny any any** rule is in place. You'll need to allow such multicast traffic in your policies explicitly. Also, remember that **Layer 2 devices** where traffic policy enforcement occurs still require **IP routing to be enabled** to function correctly with MSS.

5.3.3 Layer 2 Devices

L2 devices where traffic policy enforcement is applied will still need to enable IP routing.

5.3.4 Network Address Translation

NAT and Policy Rule Limitations

If you enable **NAT** on 7050S, 720S, or 722S platforms, be aware that the number of **supported traffic policy rules will be reduced**. This reduction is because NAT and traffic policies share the same Ternary Content Addressable Memory (TCAM) resources. Careful planning of your NAT implementation is necessary to avoid impacting your segmentation policies.

5.3.5 Access Control Lists and Policy-Based Routing

Policy Overlap and Precedence

The interaction between different policy types, specifically **Access Control Lists (ACLs)**, **Policy-Based Routing (PBR)**, and **MSS Traffic Policy rules**, requires careful consideration. If these policies are configured to apply to overlapping flow attributes, their combined effect might not be as intended. For instance, ACLs can be safely applied to **Self-IP traffic**, but keep in mind that MSS Studio isn't aware of any ACLs that impact data packets directly. You'll need to manually account for how these different policy mechanisms might interact to avoid unintended traffic behavior or security gaps.

References

6.1 Related Documents

The following documentation is available for MSS:

- 7200S Series ZTX MSS Appliance Quick Start Guide - [7200S Series ZTX MSS Appliance QSG](#)
- CloudVision Onboard Devices - [CV Onboard Devices](#)
- CloudVision Static Configuration Studio - [CV Static Configuration Studio](#)
- CloudVision MSS - [CV MSS](#)
- EOS System Configuration Guide - [EOS System Configuration Guide](#)
- NIST Zero Trust Architecture - [NIST Zero Trust Architecture](#)
- Arista MSS Technical White Paper - [Arista MSS Technical White Paper](#)
- Arista MSS Datasheet - [Arista MSS Datasheet](#)
- Arista ZTX-7250S Traffic Mapper Datasheet (ZTX Monitor Node) - [Arista ZTX-7250S Traffic Mapper Datasheet](#)
- Arista CloudVision Help Center - [Arista CloudVision Help Center](#)
- IEEE 802.1Q - Standard for Local and Metropolitan Area Networks - Bridges and Bridged Networks - [IEEE 802.1Q - Standard for Local and Metropolitan Area Networks - Bridges and Bridged Networks](#)