

Juniper Routing Director 2.5.0 User Guide

Published
2025-07-21

RELEASE
2.5.0

Juniper Networks, Inc.
1133 Innovation Way
Sunnyvale, California 94089
USA
408-745-2000
www.juniper.net

Juniper Networks, the Juniper Networks logo, Juniper, and Junos are registered trademarks of Juniper Networks, Inc. in the United States and other countries. All other trademarks, service marks, registered marks, or registered service marks are the property of their respective owners.

Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

Juniper Routing Director 2.5.0 User Guide

2.5.0

Copyright © 2025 Juniper Networks, Inc. All rights reserved.

The information in this document is current as of the date on the title page.

YEAR 2000 NOTICE

Juniper Networks hardware and software products are Year 2000 compliant. Junos OS has no known time-related limitations through the year 2038. However, the NTP application is known to have some difficulty in the year 2036.

END USER LICENSE AGREEMENT

The Juniper Networks product that is the subject of this technical documentation consists of (or is intended for use with) Juniper Networks software. Use of such software is subject to the terms and conditions of the End User License Agreement ("EULA") posted at <https://support.juniper.net/support/eula/>. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

Table of Contents

1

[About This Guide | xxv](#)

Introduction

[Overview | 2](#)

[About the Routing Director Documentation | 2](#)

[Routing Director Overview | 3](#)

[Licensing Overview | 5](#)

[GUI Overview | 6](#)

[GUI Menu Overview | 24](#)

[Personas Overview | 30](#)

[Retention Policies in Paragon Automation | 34](#)

Access and Manage Routing Director Account | 38

[Access the Routing Director GUI | 38](#)

[User Activation and Login | 39](#)

[Log into Routing Director for the First Time After Installation | 40](#)

[Log in as a New User Without an Invite \(when SMTP is not Configured\) | 40](#)

[Log in as a New User for the First Time Without an Invite \(when SMTP is Configured\) | 41](#)

[Log in as a New User With an Invite \(when SMTP is Configured\) | 41](#)

[Access a New Organization \(when SMTP is Configured\) | 42](#)

[Reset Your Password | 42](#)

[Reset your Password \(When SMTP is Configured\) | 43](#)

[Reset your Password | 43](#)

[Reset Password \(When SMTP is Not Configured\) | 44](#)

[Reset the Password of a Superuser, in Paragon Shell, as a System Administrator | 44](#)

[Reset the Password of a User from the Web GUI | 45](#)

[Reset the Password as a User Logging in For the First Time | 45](#)

LLM Connector | 47

[LLM Connector Overview | 47](#)

Configure LLM Connector | 50

Administration

Introduction | 53

Administration Overview | 53

Administration Workflow | 55

Organization Management | 57

Organization and Sites Overview | 57

Add an Organization | 58

Delete an Organization | 59

Manage Organization Settings | 60

Authentication Methods Overview | 65

Manage Identity Providers | 68

Add an Identity Provider | 71

Edit an Identity Provider | 71

Delete an Identity Provider | 71

Manage API Tokens | 72

Add an API Token | 73

Edit an API Token | 74

Delete an API Token | 74

Manage Roles | 75

Add a Role | 76

Edit a Role | 77

Delete a Role | 77

Configure SNMP | 77

Configure Webhooks to Receive Event Notifications | 81

Manage RADIUS Server Configurations | 84

Add a RADIUS Server | 85

Edit a RADIUS Server | 86

Site Management | 88

About the Sites Page | 88

Add Sites | 90

Edit and Delete Sites | 91

 Edit Site Information | 92

 Delete a Site | 92

User Management | 93

About the Users Page | 93

Predefined User Roles Overview | 97

Configure SMTP Settings in Paragon Shell | 101

 Enable SMTP Based User Management | 102

 Disable SMTP Based User Management | 105

Add Users to an Organization Overview | 107

Invite Users | 109

 Invite Users when SMTP is Enabled | 110

 Invite Users When SMTP is Disabled and Local Authentication is Enabled | 110

Manage Users and Invites | 112

 Edit User Information when SMTP is Configured | 113

 Reinvite a User | 113

 Cancel an Invitation | 114

 Revoke a User | 115

 Edit User Information when SMTP is not Configured | 115

 Reset Your Password | 116

 Delete a User | 116

Manage Paragon Shell Users | 117

 Create a User | 117

 Modify User Information | 120

 Delete a User | 122

 Logging in to Paragon Shell as a New User | 123

 Retrieve User Information from a Recovered Node in the Routing Director Cluster | 124

Manage Your Routing Director Account | 125

 Change Account Information and Password | 126

- Enable Two-Factor Authentication | 126
- Enable E-Mail Notifications When SMTP is Configured | 127
- Delete Your Routing Director Account | 127

Inventory Management | 128

About the Inventory Page | 128

Assign a Device to a Site | 133

Release a Device | 133

Replace a Device | 135

Audit Logs | 137

Audit Logs Overview | 137

About the Audit Logs Page | 138

Tags | 141

About the Tags Page | 141

Manage Tags | 143

- Add a Tag to Routing Director | 143
- Edit a Tag | 144
- Delete Tags | 144

Attribute-Based Access Control | 145

Access Control Profile Overview | 145

About the Access Control Profiles Page | 151

Add Access Control Profiles | 154

Add Access Control Profile Policies | 155

Edit and Delete Access Control Profiles | 165

- Edit an Access Control Profile | 165
- Delete an Access Control Profile | 166

Device Life Cycle Management

Introduction | 168

Device Life-Cycle Management Overview | 168

Device Onboarding Overview | 171

Device Onboarding Workflow | 173

Day-Wise Activities for Device Life Cycle Management | 177

Add Network Resource Pools and Profiles (Day -2 Activities) | 177

Prepare for Device Onboarding (Day -1 Activities) | 178

Install and Onboard the Device (Day 0 Activities) | 179

Add a Device to Routing Director | 188

- Adopt a Juniper Device | 189

- Adopt a Juniper Device Using ZTP | 190

- Add a Non-Juniper Device | 194

Move a Device to Production (Day 1 and Day 2 Activities) | 197

Field Technician User Interface | 199

Field Technician UI Overview | 199

Working with Field Technician UI Pages | 200

- Onboard a Device Page | 201

- Device List Page | 201

Onboarding Profiles | 202

Device and Interface Profiles Overview | 202

About the Device and Interface Profiles Page | 205

Add Labels | 207

Add a Device Profile | 208

Add an Interface Profile | 222

Add a Port Profile | 226

Edit and Delete a Label or Profile | 227

- Edit a Label or Profile | 228

- Delete a Label or a Profile | 229

Plan Device Onboarding | 230

Network Implementation Plan Overview | 230

About the Network Implementation Plan Page | 234

Add a Network Implementation Plan	241
Publish a Network Implementation Plan	258
Offboard a Network Implementation Plan	259
Edit a Network Implementation Plan	260
Update the Network Implementation Plan After Routing Director Upgrade	261
View Device Onboarding 	264
About the Put Devices into Service Page	264
Move a Device to Production	268
View Results of Automated Device Tests	268
Identity and Location Data of a Device	270
Remote Management Data and Test Results	272
Hardware Data and Test Results	278
Overview	278
Hardware Details for <i>Device-Name</i> Page	282
Interfaces Data and Test Results	287
Overview	288
Pluggables Details for <i>Device-Name</i> Page	292
Input Traffic Details for <i>Device-Name</i> Page	295
Output Traffic Details for <i>Device-Name</i> Page	299
Interfaces Details for <i>Device-Name</i> Page	303
Software Data and Test Results	306
Configuration Data and Test Results	308
Routing and MPLS Data and Test Results	310
Overview	311
BGP Routing Details for <i>Device-Name</i> Page	315
IS-IS Routing Details for <i>Device-Name</i> Page	319
OSPF Routing Details for <i>Device-Name</i> Page	322
RSVP Routing Details for <i>Device-Name</i> Page	325
LDP Routing Details for <i>Device-Name</i> Page	328
LSP Routing Details for <i>Device-Name</i> Page	330

- RIB Details | 333
- FIB Details | 334
- Traffic Loss | 336
- Predictor Events | 338

Device Connectivity Data and Tests Results | 343

- Connectivity Accordion | 344
- Connectivity Details Page | 347
- View Connectivity Test Results | 350

Device Management | 355

Device Management Workflow | 355

Device Licenses Overview | 357

About the Features Tab | 358

About the Licenses Tab | 360

Manage Device Licenses | 361

- Add a Device License | 361
- Delete a Device License | 362

About the Software Images Page | 363

Upload a Software Image | 366

Delete a Software Image | 371

About the Configuration Backups Page | 372

Configuration Templates Overview | 375

About the Configuration Templates Page | 376

Add a Configuration Template | 379

Edit and Delete a Configuration Template | 388

- Edit a Configuration Template | 388
- Delete a Configuration Template | 388

Preview a Configuration Template | 389

Deploy a Configuration Template to a Device | 390

Observability

Introduction | 393

Observability Overview | 393

Health Dashboard | 398

About the WAN Health Tab | 399

About the Devices Health Page | 402

About the Interfaces Health Page | 406

About the Routing Health Page | 411

About the Trust Tab | 416

About the Vulnerability Health Page | 419

About the Compliance Health Page | 423

About the Integrity Health Page | 427

Update Trust KPIs in an Air-Gapped Installation | 430

 | Update Trust KPIs using Paragon Shell | 431

About the Active Assurance Tab | 434

About the Applications Page | 435

About the L3 Network Page | 439

About the L2 Network Page | 442

About the TV & Multicast Page | 446

Troubleshoot Devices | 450

Troubleshoot Using Alerts and Alarms | 450

About the Troubleshoot Devices Page | 454

Upgrade Images on Devices | 460

About the *Device-Name* Page | 461

About the Chassis Tab | 464

About the Interfaces Tab | 466

About the Events Page | 470

Alerts Tab | 471

Alarms Tab | 474

Device Logs Tab | 477

Manage Event Templates | 480

Create an Event Template | 481

Edit Event Template Configuration | 485

Clone an Event Template | 486

Delete an Event Template | 486

Monitor Devices Using AI/ML | 488

Automatically Monitor Device and Interface Health and Detect Anomalies | 488

Device and Interface Health Monitoring and Anomaly Detection Overview | 488

View Device and Interface Health Anomalies on the GUI | 491

Detect Blackholes | 495

View Network Topology | 498

Network Topology Visualization Overview | 498

Network Visualization Options | 501

View Network Topology Details | 504

Topology Map | 504

Topology Menu Bar | 509

Network Table Overview | 515

About the Device Tab | 516

About the Link Tab | 520

About the Site Tab | 524

About the SRLGs Page | 526

About the Tunnels Tab | 528

Add a Tunnel | 535

Add Diverse Tunnels | 542

About the FlexAlgo Tab | 548

Reroute LSPs | 550

Add and Remove LSP Delegation | 553

 Add LSP Delegation | 554

 Remove LSP Delegation | 554

Dynamic Topology Workflow | 555

Configure a PCE Server | 556

Configuring PCEP on a PE Router | 557

Configuring Topology Acquisition Using BGP-LS | 559

Enable PCEP Security | 560

Routing Observability | 565

Routing Observability Overview | 565

Enable Routing Observability and View Routing Analytics Data | 566

About the Route Explorer Page | 566

About the Routing Status Tab | 567

About the Routing Updates Tab | 571

About the Adjacencies Tab (Route Explorer) | 573

Rules | 576

Rules Overview | 576

 Understand Custom Rules | 577

About the Custom KPI Collection Page | 578

About the Rule Instantiations Tab | 578

About the Custom Rules Tab | 580

 About the Rules List Tab | 581

 About the Rule Helper Files Tab | 582

About the Alerts Tab (Custom Rules) | 584

About the Monitor Instance-Name Page | 586

 Alerts Tab | 586

 KPI Tab | 588

Manage Custom Rules | 590

- Upload a Custom Rule | 590
- Download a Custom Rule | 591
- Add Tags to a Custom Rule | 591
- Delete a Custom Rule | 592

Manage Rule Helper Files | 593

- Upload a Rule Helper File | 593
- Download a Rule Helper File | 594
- Delete a Rule Helper File | 594

Add a Rule Instance | 594**Manage Rule Instances | 596**

- Edit a Rule Instance | 596
- Add Tags to a Rule Instance | 597
- Delete a Rule Instance | 598

View KPI Performance and Alerts (Custom Rules) | 598

- Overview | 598
- Components on the KPI Tab | 599
- KPI Performance Graphs | 599

5

Trust and Compliance**Introduction | 603****Trust and Compliance Overview | 603****Perform Compliance Scan and Manage Checklists | 604****Manage Trust Settings and Trust Scores | 607****Compliance Standards Overview | 607****About the Compliance Benchmarks Page | 608****About the Compliance Tailorings Page | 610****Example: Create a Tailoring Document for NTP Settings | 612****About the Compliance Checklist Page | 613****Add a Checklist Template | 615****Add Checklist for a Device | 616**

Import Scans and Update Rule Results in a Checklist | 617

Trust Plans Overview | 617

About the Network Score Formula Page | 619

Trust Score Overview | 621

Manage Compliance Scans | 623

Compliance Scans Overview | 623

About the Compliance Page | 624

Perform Custom Compliance Scans | 626

Analyze Scan Results | 630

Manage Vulnerabilities | 631

Vulnerabilities Overview | 631

About the SIRT Advisories Page | 632

About the Proactive Bug Notifications Page | 635

Monitor Integrity | 638

Integrity of the Hardware and Software on the Network | 638

About the Software End of Life Page | 639

About the Hardware End of Life Page | 641

6

Service Orchestration

Introduction | 646

Service Orchestration Overview | 646

Service Provisioning Workflow | 648

View Service Design Catalog | 651

Service Design Overview | 651

About the Service Designs Page | 653

Manage Service Design Versions | 657

Set Service Design Default Version | 657

Install Latest Available Service Design Version | 657

Upgrade Service Design Version for a Service Instance	658
Upload a Customized Service Design	659
Manage Customers 	664
About the Customer Inventory Page	664
Add a Customer	666
Edit and Delete Customers	667
Edit Customer Details	667
Delete a Customer	668
Add Resources for Network Services 	669
Resource Instance Overview	669
Add Network Resources for Provisioning Services	670
About the Resource Instances Page	672
Add a Resource Instance	674
Add Network Resource Pools	676
Upload Network Resource Pools for Device Onboarding in JSON Format	677
Add Network Resource Pools by Using REST APIs	677
Sample Files	680
Configure Topology Resources in Network Implementation Plan	743
Configure Resource Pools for Resource Instances	745
Configure a Network Resource Pool	745
Create a Routing Resource Pool	746
Create a Layer 3 Resource Pool	748
Create a Layer 2 Resource Pool	750
Create a Topology Resource Pool	752
Create a VPN Resource Pool	763
Modify, Delete, and Restore Resource Instances	766
Modify a Resource Instance	766
Delete a Resource Instance	767
Restore a Resource Instance	767
Manage Placement Configurations for Service Instances	768

- Manage Placement Configurations by Using the Routing Director GUI | 768
- Manage Placement Configurations by Using the Service Orchestration cMGD CLI | 770

View Network Resources | 771

Provision VPN Services | 773

Add an L3VPN Service Instance | 773

Add L3VPN Service Site Details | 777

- Add L3VPN Site | 777
- Add L3VPN Site Network Access Parameters | 779
- Add Access Diversity Parameters for L3VPN Site Network Access | 786
- Add Routing Protocol for L3VPN Site Network Access | 788

Add an EVPN or EVPN-VPWS Service Instance | 791

Add EVPN or EVPN-VPWS Service Site Details | 796

- Add EVPN or EVPN-VPWS Site | 796
- Add EVPN or EVPN-VPWS Site Network Access Parameters | 799
- Add Access Diversity Parameters for EVPN and EVPN-VPWS Site Network Access | 804
- Add Connection Parameters for EVPN or EVPN-VPWS Site Network Access | 806

Add an L2 Circuit Service Instance | 810

Add L2 Circuit VPN Nodes | 814

Modify a Service Instance | 822

Manage Service Instances | 828

Service Instance Overview | 828

About the Service Instances Page | 829

View Service Instance Details | 836

View Passive Assurance Monitoring Data | 843

View EVPN or EVPN-VPWS Passive Assurance Monitoring Results | 845

- View L2VPN or EVPN-VPWS Health Check Results | 846
- View OAM Health Check Results | 848
- View LLDP Health Check Results | 850

View L2 Circuit Passive Assurance Monitoring Data | 852

- View L2 Circuit Service Health Check Results | 852
- View Health Check Results for L2 Circuit Connection Status | 854

View L3VPN Passive Assurance Monitoring Results | 854

- View BGP Health Check Results | 856
- View OSPF Health Check Results | 858
- View L3VPN Service Health Check Results | 862

Physical Interfaces Health Check Results | 864

- View Health Check Results for Physical Interface Parameters | 866

Logical Interfaces Health Check Results | 869

- View Health Check Results for Logical Interface Operation Status | 870

Performance Graphs to View Passive Assurance Monitoring Data | 871

Monitor Service Order Execution Workflows | 874

About the Service Orders Page | 874

About the Workflows Page | 880

About the Workflow Runs Page | 882

View Workflow Run Details | 885

7

Active Assurance

Introduction | 890

Active Assurance Overview | 890

Active Assurance Workflow | 891

Active Assurance Terminologies | 892

Test Agents | 897

Test Agents Overview | 897

About the Test Agents Page | 900

Add a Test Agent | 908

About the *Test Agent-Name* Page | 911

Install Test Agent Application | 929

- Deploy a Test Agent Using Container Image | 930

- Build a Test Agent Container Image | 931
- Install a Test Agent as a Native Application in Linux | 932

Install Test Agent Appliance | 934

Tests and Monitors | 944

Tests and Monitors Overview | 944

About the Measurement Designer Page | 955

Create a Test | 958

Create Input Variables | 1097

About the Tests Page | 1100

About the *Test-Name* Page | 1104

Create a Monitor | 1110

About the Monitors Page | 1247

About the *Monitor-Name* Page | 1253

View Stream Details | 1258

Test Templates | 1288

About the Test Templates Page | 1288

Create a Test Template | 1291

8

Network Optimization

Introduction | 1294

Network Optimization Overview | 1294

Network Optimization Intents | 1296

Path Intent Workflow | 1296

About the Network Optimization Intent Page | 1298

About the Path Intent Tab | 1300

Add a Path Intent | 1302

About the Tunnel Profiles Tab | 1304

Add a Tunnel Profile | 1305

About the Optimization Profiles | 1310

Add an Optimization Profile | 1311

About the Endpoint Profiles | 1313

Add an Endpoint Profile | 1314

About the Device Tab (Network Optimization) | 1316

About the Link Tab (Network Optimization) | 1320

About the Tunnels Tab (Network Optimization) | 1324

Network Planner

Introduction | 1333

Planner Overview | 1333

Planner Workflow | 1334

Offline Models | 1336

About the Working Models Page | 1336

Create an Offline Network | 1338

Exhaustive Failure Simulation | 1339

View Reports | 1340

Offline Topology Map | 1343

About the Offline Topology Page | 1343

About the Device Tab (Planner) | 1345

Create a Node (Planner) | 1347

About the Link Tab (Planner) | 1348

Create a Link (Planner) | 1352

About the Tunnel Tab (Planner) | 1355

Create a Tunnel (Planner) | 1357

About the Demands Tab (Planner) | 1360

Create a Demand (Planner) | 1362

About the Interface Tab (Planner) | 1363

Create an Interface (Planner) | 1365

Paragon Shell CLI Reference

Introduction | 1368

Paragon Shell Overview | 1368

A Quick Tour of Paragon Shell | 1369

Operational Mode Commands | 1376

file copy | 1377

monitor | 1379

request paragon backup | 1381

request paragon cluster pods reset | 1389

request paragon cluster upgrade | 1391

request paragon config | 1398

request paragon deploy | 1400

request paragon deploy cluster | 1403

request paragon destroy cluster input | 1406

request paragon fix-permission | 1408

request paragon health-check | 1409

request paragon load | 1412

request paragon repair-node | 1413

request paragon replace-node | 1414

request paragon restore | 1416

request paragon reboot | 1418

request paragon running-config | 1420

request paragon shutdown | 1422

request paragon service | 1424

request paragon ssh | 1426

request paragon ssh-key | 1427

request paragon storage cleanup | 1431

request paragon super-user password reset | 1433

request system decrypt password | 1435

request system reboot | 1436

request paragon trust data update hostname | 1438

show configuration paragon cluster | 1439

show configuration paragon cluster applications (o) | 1448

show configuration paragon cluster common-services (o) | 1451

show host disk usage | 1454

show paragon backup | 1457

show paragon certificate expiry-date certificate-type | 1462

show paragon cluster (o) | 1465

show paragon cluster details | 1467

show paragon cluster namespaces | 1468

show paragon cluster nodes | 1470

show paragon cluster pods | 1473

show paragon cluster pods namespace healthbot sort | 1476

show paragon images version | 1479

show paragon images version namespace | 1481

show paragon pvc details | 1484

show paragon version | 1489

Configuration Mode Commands | 1491

delete paragon cluster | 1491

load set | 1497

set paragon cluster applications | 1498

set paragon cluster common-services | 1504

set paragon cluster insights victoria-metrics | 1507

set paragon cluster install | 1509

set paragon cluster kubernetes | 1511

set paragon cluster mail-server | 1512

set paragon cluster nodes | 1515

set paragon cluster ntp | 1516

set paragon cluster papi | 1518

set paragon monitoring | 1519

set system login | 1522

show paragon cluster (c) | 1524

show paragon cluster applications (c) | 1529

show paragon cluster common-services (c) | 1532

Troubleshooting Commands | 1536

Troubleshoot Using the Paragon Shell CLI Commands | 1536

- Overview | 1537
- request paragon support information | 1537
- request paragon troubleshooting information | 1538
- Other Troubleshooting Commands to Debug Issues | 1541

request paragon support information | 1543

request paragon troubleshooting information | 1546

request paragon debug | 1550

request paragon debug get-tsdb-data | 1552

request paragon debug insights-kafka-data | 1554

request paragon debug kafka | 1556

request paragon debug logs | 1558

request paragon debug logs namespace | **1561**

request paragon debug logs namespace healthbot | **1561**

request paragon debug logs namespace foghorn | **1563**

request paragon debug logs namespace papi | **1565**

request paragon debug logs namespace northstar | **1567**

request paragon debug logs namespace airflow | **1569**

request paragon debug postgres | **1571**

request paragon debug redis | **1573**

Service Orchestration | 1575

About the Service Orchestration cMGD CLI | **1576**

Access the Service Orchestration cMGD CLI | **1577**

Directories in the Service Orchestration cMGD | **1578**

Service Orchestration cMGD CLI Commands | **1578**

set foghorn:core org-id | **1583**

set service design default version | **1584**

set dag state failed | **1586**

show service order network consumers | **1587**

show service order network resources | **1589**

show service order status | **1591**

show service order as-json | **1594**

show service order as-yaml | **1597**

show service designs | **1601**

show device dependent configuration | **1604**

show insights configuration | **1607**

show configuration foghorn:customers | **1610**

show service instance lock status | **1611**

show service instance lock token | **1614**

show service catalog projects | **1616**

show service projects | **1618**

show service order history | **1619**

show airflow dags | **1621**

show airflow dag-runs | **1623**

show paragon orgs | **1626**

request service project add | **1627**

request service catalog project add | **1629**

request service order sync | **1630**

request network resources load | **1631**

request service order load | **1633**

request service order place | **1634**

request service order modify | **1636**

request service catalog project | **1637**

request service catalog design delete | **1639**

request service order delete | **1640**

request service order submit | **1641**

request service order provision | **1642**

request service design install | **1643**

request service design uninstall | **1644**

request create graphdb | **1646**

request service order set | **1647**

request service instance modify design-version | **1648**

request service instance lock | **1649**

request service instance unlock | **1651**

About This Guide

Juniper® Routing Director provides end-to-end transport network automation and simplifies the adoption of network automation for device, network, and service life cycles from Day 0 to Day 2. Routing Director enables network operations teams to improve productivity and operational efficiency by eliminating manual tasks, processes, and workflows that are often repetitive and prone to human error.

Routing Director provides the following WAN automation use cases for service providers and enterprises:

- Device life-cycle management
- Network trust and compliance
- Observability
- Service orchestration
- Active assurance
- Network optimization
- Network planner

Use this guide to understand the various use cases in Routing Director. This guide provides overviews, workflows, and procedures that help you understand the use cases and perform various tasks in Routing Director.

1

PART

Introduction

- Overview | 2
 - Access and Manage Routing Director Account | 38
 - LLM Connector | 47
-

CHAPTER 1

Overview

IN THIS CHAPTER

- [About the Routing Director Documentation | 2](#)
- [Routing Director Overview | 3](#)
- [Licensing Overview | 5](#)
- [GUI Overview | 6](#)
- [GUI Menu Overview | 24](#)
- [Personas Overview | 30](#)
- [Retention Policies in Paragon Automation | 34](#)

About the Routing Director Documentation

This Routing Director user guide provides overviews, workflows, and procedures that help you understand the use cases and perform various tasks in Routing Director.

This guide is available on the [Juniper Routing Director Documentation](#) page, which contains links to the documentation for all Routing Director releases. Select the appropriate release from the release picker drop-down at the top right of the page to view documentation for the current release. [Table 1 on page 2](#) lists the documentation that is available on the page.

Table 1: List of Routing Director Documents

Documentation	Location
Release Notes	PLAN section of the Juniper Routing Director Documentation page
Licensing Guide	PLAN section of the Juniper Routing Director Documentation page
Installation and Upgrade Guide	SET UP section of the Juniper Routing Director Documentation page

Table 1: List of Routing Director Documents *(Continued)*

Documentation	Location
Quick Start: Routing Director	SET UP section of the Juniper Routing Director Documentation page
Quick Start: Onboard Juniper Networks Devices to Paragon	SET UP section of the Juniper Routing Director Documentation page
Monitoring and Troubleshooting Guide	MANAGE section of the Juniper Routing Director Documentation page

For purchasing a license, contact your [Juniper Networks](#) sales representative.

For downloading the Routing Director software, click [here](#).

For receiving technical support from the Juniper Networks Technical Assistance Center (JTAC), click [here](#).

Routing Director Overview

IN THIS SECTION

- [Benefits](#) | 5

Service providers, cloud providers, and enterprises are facing an increase in the volume, velocity, and types of traffic. This creates both unique challenges (increased user expectations and expanded security threats) and fresh opportunities (new generation of 5G, IoT, distributed edge services) for network operators.

To accommodate rapid changes in traffic patterns, service providers and enterprises need to quickly detect and troubleshoot devices and service issues, and make changes to service configurations in real-time. Any misconfiguration due to human errors can lead to service outages. Investigating and resolving these issues can be a time-consuming process.

Juniper® Routing Director is a WAN automation solution that enables service provider and enterprise networks to meet these challenges. Juniper's solution delivers an experience-first and automation-driven network that provides a high-quality experience to network operators.

Routing Director is based on a modern microservices architecture with open APIs. Routing Director is designed with an easy to use UI that provides a superior operational and user experience. For example, Routing Director implements different persona profiles (such as network architect, network planner, field technician, and Network Operations Center [NOC] engineer) to enable operators to understand the different activities in the device life-cycle management (LCM) process. For details, see ["Personas Overview" on page 30](#).

Routing Director takes a use case-based approach to network operations. When you execute a use case, Routing Director invokes all the required capabilities of that use case, runs a workflow (if necessary) and presents you with a completed set of tasks that implements the use case.

Routing Director supports the following use cases (explained at a high-level):

- **Device life-cycle management (LCM)**—Allows you to onboard, provision, and then manage a device. Routing Director automates the device onboarding experience, from shipment through service provisioning, thus enabling the device to be ready to accept production traffic. For more information, see ["Device Life-Cycle Management Overview" on page 168](#).
- **Observability**—Allows you to visualize the network topology, provision tunnels, and monitor the devices and the network. You can also view the device and network health and drill down into the details. In addition, Routing Director notifies you about network issues using alerts, alarms, and events, which you can use to troubleshoot issues affecting your network. You to actively monitor the overall routing health in real time using the route explorer dashboard and the route topology map. For more information, see ["Observability Overview" on page 393](#).
- **Trust and compliance**—Automatically checks whether the device complies with the rules defined in the Center for Internet Security (CIS) benchmarks document. In addition, Routing Director also checks the configuration, integrity, and performance of the device and then generates a trust score that determines the device's trustworthiness. For more information, see ["Trust and Compliance Overview" on page 603](#).
- **Service Orchestration**—Enables you to streamline and optimize the delivery of network services and thereby improving efficiency and reducing the risk of errors. A service can be any point-to-point, point-to-multipoint or multipoint-to-multipoint connection. For example, Layer 3 VPNs or EVPNs. For more information, see ["Service Orchestration Overview" on page 646](#).
- **Active Assurance**—Enables you to actively monitor and test the network's data plane by generating synthetic traffic using Test Agents. Test Agents are measurement points deployed in certain routers in your network. In addition, there are pre-deployed Test Agents in certain clouds, such as AWS. These Test Agents are capable of generating, receiving, and analyzing network traffic and therefore enables you to continuously view and monitor both real-time and aggregated result metrics. For more information, see ["Active Assurance Overview" on page 890](#).

- **Network Optimization**—Enables you to optimize the utilization of network resources, enhance network performance, and ensure reliable and efficient delivery of data across the network. This is achieved by managing the life-cycle of Label-Switched Paths (LSPs), through an intent-based approach. For more information, see "[Network Optimization Overview](#)" on page 1294.
- **Network Planner**—Provides in-depth network views and reports on how the network is performing in a particular failure scenario, all without impacting your production environment. For more information, see "[Planner Overview](#)" on page 1333.

For details about these use cases and other features of Routing Director, refer to the corresponding sections in the Routing Director User Guide.

Benefits

- Automate the onboarding and provisioning of devices
- Simplify and accelerate service delivery
- Reduce manual effort and timelines by using automation
- Intent-based service orchestration solution with built-in active assurance
- Can be deployed in an air-gapped environment and in a private network thereby preventing security-related risks.

RELATED DOCUMENTATION

[Access the Routing Director GUI](#) | 38

[GUI Menu Overview](#) | 24

Licensing Overview

To use Routing Director and its features, you need a:

- **Product Entitlement**—To use Routing Director and its use cases.



NOTE: Product entitlements are honor-based and not enforced for Routing Director Release 2.5.0.

- **Device License**—To use the features on a device that you onboarded.

For more information on how to add a device license in Routing Director, see "[Device Licenses Overview](#)" on page 357.

To purchase a product entitlement or a device license, you can contact your [Juniper Sales Representative or Business Partner](#). After you complete your purchase, you can download the license file and manage the license by using the [Juniper Agile Licensing](#) (JAL) portal. You can also choose to receive the license file over an email.

RELATED DOCUMENTATION

| [Juniper Agile Licensing Overview](#)

GUI Overview

IN THIS SECTION

- [Menu and Banner | 7](#)
- [Breadcrumbs and GUI Elements in Landing Pages | 11](#)
- [Sort, Resize, Filter, and Search Icons, and Related GUI Elements | 12](#)
- [Page Display, Navigation, and Related GUI Elements | 16](#)
- [View, Add, and Remove Favorite Pages | 17](#)
- [Filter Data in a Table | 18](#)

The Routing Director GUI provides an easy to use, single pane of glass experience that allows you to access the different use cases and features.

To access the Routing Director GUI, you must log in using your Routing Director account. For more information, see "[Access the Routing Director GUI](#)" on page 38. After you log in successfully to the Routing Director GUI, you are taken to the Troubleshoot Devices page, which displays the devices belonging to your organization and enables you to manage the devices. For more information, see "[About the Troubleshoot Devices Page](#)" on page 454.

In this topic, we'll discuss some commonly used elements and features of the Routing Director GUI.

Menu and Banner


The two elements of the Routing Director GUI that you'll use frequently are as follows:

- **Menu:** The menu, which is available at the left-side of the GUI, is minimized by default. You can hover over or click inside the menu to expand the menu. A sample of the expanded menu is shown in [Figure 1 on page 7](#).

You can expand the menu and click different menu entries to navigate to the different pages in the Routing Director. For details about the menu, see ["GUI Menu Overview" on page 24](#).

- **Banner:** The banner, which is displayed at the top of the page (see [Figure 1 on page 7](#)) contains several icons and GUI elements that you're likely to use regularly. These icons and GUI elements are explained in [Table 2 on page 8](#).

Figure 1: Sample Page Showing Menu and Banner



1– Menu toggle icon	6– Organization drop-down
2– Expand/Collapse Navigation	7– Settings Menu
3– Menu bar and expanded menu	8– Help (?) icon
4– Banner	9– User account icon
5– Organization name	

Table 2: Banner Icons and GUI Elements

Description	Function
Menu Toggle	Click either the menu toggle icon (the icon with three horizontal bars) in the top left of the banner or the Expand/Collapse Navigation button (> or <) to toggle the visibility of the Routing Director menu. If the menu was previously hidden, it is displayed, and the menu is hidden if it was previously displayed.
Organization drop-down	The Organization drop-down displays the current organization that you are accessing. Click the Down arrow next to the organization name expand the drop-down. You can: • View the list of organizations to which you have access. You can click an organization name to switch context to that organization. • Click Create Organization to add an organization. For more information, see "Add an Organization" on page 58.

Table 2: Banner Icons and GUI Elements *(Continued)*

Description	Function
Settings Menu	Click the gear icon to quickly access the following settings from the banner instead: • System Settings—View, configure, and manage the settings of an organization. For more information, see "Manage Organization Settings" on page 60. • Users—View details of the existing users and the users invited to access the organization. For more information, see "About the Users Page" on page 93. • Audit Logs—View audit logs generated by user-initiated tasks based on the time interval you select. For more information, see "About the Audit Logs Page" on page 138. • Control Profiles—View access control profiles, which provides granular access to Routing Director resources (for example, devices, service instances, network implementation plan, and so on) and manage tags. For more information, see "About the Access Control Profiles Page" on page 151.


Table 2: Banner Icons and GUI Elements *(Continued)*

Description	Function
Help (?) menu	Click the (?) (help) icon to access the help menu. The Need help? page is displayed, which provides links to the following: • Release Notes—Opens the Release Notes panel within the application, which lists the new and changed features and the bug fixes in the current software release. • Quick Help—Opens the Quick Help panel within the application, which contains the topics that explain how to use Routing Director. You can use the Featured tab to access featured topics or the All Topics tab to access all topics. • API Docs—Opens the REST API documentation. Use this guide to view the list of Routing Director REST APIs. • Juniper Support—Opens the Juniper Support portal. Use this portal to create and manage your support cases. • About—Opens the About panel, which provides information about the software release and copyright information. The About link is available in the bottom-right corner of the Need help? page.
User account icon	Click the user account icon to access the user account menu. This menu displays your name and e-mail address, and you can do the following: • Manage your account: Click My Account to open the My Account page, where you can modify your account, password, and other information. See "Manage Your Routing Director Account" on page 125. • Log out of Routing Director: Click Logout to log out of the GUI. You are logged out and taken to the Juniper login page.

Breadcrumbs and GUI Elements in Landing Pages

Figure 2 on page 11 shows the breadcrumbs, page help, and other GUI elements or icons, and Table 3 on page 11 provides a high-level explanation of their functions.

Figure 2: Sample Page Showing Breadcrumbs, Page Help Icon, and Other GUI Elements



1– Breadcrumbs	5– Edit icon
2– Page Help icon	6– Delete icon
3– More drop-down	7– Favorite icon
4– Add or Create icon	

Table 3: Breadcrumbs, Page Help Icon, and Other GUI Elements or Icons

Description	Function
Breadcrumbs	The breadcrumbs in the Routing Director display the menu structure and provide an alternative way to navigate the menu. Click the Down arrow next to a breadcrumb to access the menu entries at that menu level.
Page Help icon	Click or hover over the page help (?) icon to view help text for the page and access the More... link. You can click the More... link to open the in-application help topic for that page.

Table 3: Breadcrumbs, Page Help Icon, and Other GUI Elements or Icons *(Continued)*


Description	Function
More drop-down	The More drop-down provides additional options for tasks that you can perform on a page.
Add or Create (+) icon	Used to add or create an entity; for example, create a site.
Edit (pencil) icon	Used to modify an existing entity; for example, modify a site.
Delete (trash can) icon	Used to delete an entity; for example, delete a site.
Favorite icon	Used to mark a page as a favorite page or remove a page that was previously marked as a favorite. See "View, Add, and Remove Favorite Pages" on page 17 .

Sort, Resize, Filter, and Search Icons, and Related GUI Elements

[Figure 3 on page 13](#) shows the sort, filter, search, and related GUI elements that you typically encounter on landing pages (for example, Sites). [Table 4 on page 13](#) lists these icons and provides a high-level explanation of their functions.

 **NOTE:** The search and filter icons might not be available on some pages.

Figure 3: Sample Page with Sort, Resize Columns, Filter, Search, and Related GUI Elements



1– Sort icons	4– Search icon
2– Resize column icon	5– Column and Page Preferences Menu
3– Filter icon	

Table 4: Sort, Resize, Filter, Search, and Related GUI Elements

Description	Function
Sort icons	The sort icons next to a column label in a table (grid) indicate that the data can be sorted (in ascending or descending order) based on that column. To sort the data, click the column label. The corresponding sort icon changes color to indicate whether the data is sorted in ascending or descending order.
Column Resize icon	In some tables, columns can be resized by moving your mouse between two column names until you see the column resize icon. You can then left-click your mouse, and hold and drag the mouse resize the column.
Re-arrange columns	To move a column, click inside a column label, hold and drag to move the column to where you want it to be placed, and release.

Table 4: Sort, Resize, Filter, Search, and Related GUI Elements *(Continued)*

Description	Function
Filter icon (funnel)	You can apply one or more filters to the data in the table and, if needed, save the filters. Hover over or click the filter icon to access the filtering menu. For more information, see "Filter Data in a Table" on page 18.
Search icon (magnifying glass)	You can click the search icon search the data and, if needed, save the search as a filter. • Click the Search icon and enter one or more keywords, and press Enter. The data displayed in the table is filtered based on the keywords that you entered. • To save the search as a filter so that it can be reused later, click Save. For details, see "Filter Data in a Table" on page 18. • To clear a search, click the X icon. The unfiltered data is displayed in the table.

Table 4: Sort, Resize, Filter, Search, and Related GUI Elements *(Continued)*

Description	Function
Vertical Ellipsis icon	Click or hover over the vertical ellipsis to access the column and page preferences menu. You can do the following: • Show or hide columns in the table (grid): 1. Hover over or click Show/Hide Columns to view the list of columns that you can display in the table. The check box next to the column indicates whether the column is displayed (check box is selected) or not (check box is cleared). 2. (Optional) Select the check boxes corresponding to the columns that you want to display in the table. The selected columns are displayed in the table. 3. (Optional) Clear the check boxes corresponding to the columns that you do not want to display. The cleared columns are no longer displayed in the table. • Reset the page preferences and remove any previously applied filters: 1. Hover over the vertical ellipsis menu and click Reset Preference. A message appears asking you to confirm the reset. 2. Click Yes. The page preferences are reset and the default view is displayed.

Page Display, Navigation, and Related GUI Elements

Figure 4 on page 16 shows the GUI elements related to page display and navigation, which that you typically encounter on landing pages (for example, Sites). Table 5 on page 16 lists these GUI elements and provides a high-level explanation of their functions.

Figure 4: Sample Page Showing Display, Navigation, and Related GUI Elements



1– Total number of entries (items) available	5– Page numbers
2– Refresh icon	6– Navigate to the next page icon
3– Number of items displayed per page	7– Go to (page number)
4– Navigate to the previous page icon	

Table 5: Page Display, Navigation, and Related GUI Elements

Function	Description
<i>Total-number</i> [of] items	Displays the total number of items or entries available on a page.
Refresh icon	Typically, pages in the Routing Director GUI refresh automatically. However, you can click the Refresh icon to trigger a manual refresh if needed.

Table 5: Page Display, Navigation, and Related GUI Elements (*Continued*)

Function	Description
Display options	This field displays the number of entries per page in the table (grid). You can click the number and select the number of items that you want to display.
Previous Page (<) icon	For tables displaying two or more pages, click < to go to the previous page.
Page numbers	Displays one or more page numbers depending on the number of pages of items (entries) displayed. Click the page number to go to that page.
Next Page (>) icon	For tables displaying two or more pages, click > to go to the next page.
Go to <i>page-number</i>	For tables displaying two or more pages, enter the page number in the text box and press Enter to go to that page.

View, Add, and Remove Favorite Pages

In Routing Director, you can mark pages that you frequently use as favorites, so that you can access such pages easily. You can view existing favorites in the Favorites menu, remove existing favorites, or add pages as favorites. A sample page showing the Favorites menu, icons, and so on is shown in [Figure 5 on page 18](#).



NOTE: The Favorites menu appears only if at least one page marked as a favorite.

You can do the following:


- View or access favorite pages: You can use the Favorites menu to view and access existing favorite pages.
- Add a page as a favorite: You can add a page as a favorite in one of the following ways:
 - By clicking the star icon next to the menu entry.
 - By clicking the star icon at the top right corner of a page (below the Routing Director banner).

When you add a page as a favorite, it appears under the Favorites menu. The star icon is shaded (filled), which indicates that the page is a favorite.

- Remove a page as a favorite: You can remove a page as a favorite in one of the following ways:
 - By clicking the shaded star icon in the Favorites menu.
 - By clicking the shaded star icon next to the menu entry.
 - By clicking the shaded star icon at the top right corner of a page.

When you remove a page as a favorite, it no longer appears in the Favorites menu. The star icon changes to empty (unshaded), which indicates that the page is not a favorite.

Figure 5: Sample Page with Favorites Menu, and Add, or Remove Favorite Icons



Filter Data in a Table

Routing Director enables you to filter the data displayed in a table (grid) based on filter criteria. You can specify one or more criterion, and use conditional operators (AND or OR) to create a combination of filter criteria.

Figure 6 on page 19 shows the expanded filter menu with and without filters and Figure 7 on page 19 shows a sample page on which filter criteria are applied. Table 6 on page 20 explains the different icons and GUI elements related to filters (as shown in Figure 7 on page 19).

Figure 6: Filter Menu with and without Filters




Figure 7: Sample Page Showing Filter Criteria Applied




Table 6: Icons and GUI Elements Related to Filters

Function	Description
Filter criteria field (text box)	This field (text box) displays the filter criteria that was previously specified. You can enter additional criteria by using the Add (+) icon.
Add criterion icon (+)	Click the + icon to add a filter criterion. For details, see "Add Filter Criteria" on page 20 .
Apply filter criteria icon (✓)	Click the check mark icon (✓) to apply the filter criteria that you specified. The filtered data is displayed in the table.
Clear all filters icon (X)	To clear all the applied filter criteria and display unfiltered data, click the X icon.
Filter icon (funnel) and drop-down	Hover over or click the filter icon or the down arrow button to access the menu to toggle the display of filters and access previously saved filters. See Figure 6 on page 19 .
Save filter button	To save the filter criteria so that you can reuse it later, click Save and follow the instructions in Step "5" on page 21 .

Add Filter Criteria

To add one or more filter criteria:

1. Do one of the following:

- If no filters are present, click the filter (funnel) icon and select **Show advanced filter** (see [Figure 6 on page 19](#)).
- If one or more filters are already present, click the Add (+) icon above the table (see [Figure 7 on page 19](#)).

A page appears displaying the fields related to filter criteria.

2. Configure the fields as described in [Table 7 on page 22](#).



NOTE: Fields marked with an asterisk (*) are mandatory.

3. Click **Add**.

The data in the table (grid) is filtered based on the criteria that you specified. The filter criterion appears on the top of the table (grid).

4. (Optional) Do one of the following:

- Specify additional filter criteria by selecting the **Operator** (see [Table 7 on page 22](#)) and configure the rest of the fields as explained in [Step "2" on page 20](#)).
- Click **Close** to close the pop-up.

You are returned to the previous page.

5. (Optional) To save the filter criteria so that you can reuse it later, click **Save**.

The Save Filter page appears.

- a. Enter a name for the filter in the **Name** text box.
- b. To set the filter as a default, click the **Set as default** toggle button.



NOTE:

- When you set a filter as a default, Routing Director automatically applies the filter on the page, and displays the filtered data.

c. Click **OK**.

A confirmation message appears indicating that the save operation was successful.

You can access saved filters using the funnel (filter) icon.



NOTE: Saved filters are stored in the local storage of the browser that you use to access Routing Director. If you clear your browser's local storage, the filters are cleared.

Table 7: Fields on the Add Criteria Pop-Up

Field	Description
Operator	NOTE: This field appears only when you've already entered one filter criterion and want to enter the second or subsequent criteria. Select the logical operator for the filter criterion that you are specifying: • AND: Data is filtered only when both the filter criteria are met. • OR: Data is filtered when one of the filter criteria is met.
Field	Select the field (parameter) that you want to use as a filtering criterion. For example, on the Sites page, you can select Name, Country, or Address as a filtering criteria.
Condition	Select the filtering condition that you want to use in the filter. A filtering condition can be: • A mathematical operator; for example, = (equal to) or != (not equal to). • A keyword; for example, starts with, Includes, or In.
Value	Specify one or more values (depending on the condition that you specified) on which to filter the data.

Apply a Saved Filter

To apply a previously saved filter:

1. Hover over or click the filter icon (funnel).

The Filter menu appears.

2. Click the filter that you want to apply.

The filtered data is displayed in the table.

Mark a Saved Filter as Default

To mark a previously saved filter as a default:

1. Hover over or click the filter icon (funnel).

The Filter menu appears.

2. Hover over the filter that you want to mark as a default and click the star icon that appears next to the filter's name.

The star icon is shaded (filled), which indicates that the filter is now a default. The next time that you access the page, the default filter is applied and the filtered data is displayed in the table.

Delete a Saved Filter

To delete a previously saved filter:



WARNING: When you trigger the deletion of a filter, it is deleted immediately. You cannot recover the filter. So, ensure that you check the filter that you want to delete before triggering a delete operation.

1. Hover over or click the filter icon (funnel).

The Filter menu appears.

2. Hover over the filter that you want to delete.

A delete icon (X) appears next to the filter name.

3. Click the delete (X) icon.

The filter is deleted. If the filter was previously saved as a default, then the filter is no longer applied on the page.

GUI Menu Overview

IN THIS SECTION

- [Observability Menu | 25](#)
- [Trust Menu | 26](#)
- [Orchestration Menu | 27](#)
- [Inventory Menu | 28](#)
- [Network Optimization | 29](#)
- [Planning | 30](#)

The Routing Director GUI menu enables you to access the different use cases and features. The tasks that you can perform are based on the roles and access privileges (capabilities) that you're assigned as a Routing Director user. For more information, see ["Predefined User Roles Overview" on page 97](#).

The menu bar is available on the left side of the Routing Director GUI. You can toggle the menu by using the menu icon (three horizontal lines) on the banner. You can also access the menu by using the breadcrumbs, that are displayed just below the banner, on every page. For more information, see ["GUI Overview" on page 6](#).

[Table 15 on page 31](#) shows the top-level menu items (sub-menus) in the Routing Director GUI.

Table 8: Routing Director Main Menu

Menu Entry	Description
Favorites	Displays the pages that are marked as favorites. For more information, see Table 2 on page 8 . NOTE: This menu appears only if you have at least one page marked as a favorite.
Observability	Access the tasks and features related to the observability use case. See Table 9 on page 25 .
Trust	Access the tasks and features related to the trust and compliance use case. See Table 10 on page 26 .

Table 8: Routing Director Main Menu *(Continued)*

Menu Entry	Description
Orchestration	Access the tasks and features related to the service orchestration use case. See Table 11 on page 27 .
Inventory	Access the tasks and features related device life-cycle management, Test Agents, and site. See Table 12 on page 28 .
Network Optimization	Access the tasks and features related to the network optimization use case. See Table 13 on page 29 .
Planning	Access the tasks and features related to the planner use case. See "Planning" on page 30 .
Device List	Access the field technician UI for the list of devices to be onboarded. For more information, see "Working with Field Technician UI Pages" on page 200 . NOTE: This menu entry appears only when you log in as a user with the Installer role.

Observability Menu

[Table 9 on page 25](#) displays the menu entries for the observability use case and links to relevant topics that you can refer to for more information.

Table 9: Observability Menu Entries

Menu Entry	Description
Health (sub-menu)	
Health Dashboard	See "About the Health Dashboard" on page 398 .
Troubleshoot Devices	See "About the Troubleshoot Devices Page" on page 454 .
Events	See "About the Events Page" on page 470 .
Custom KPI Collection	See "About the Custom KPI Collection Page" on page 578 .
Active Assurance (sub-menu)	

Table 9: Observability Menu Entries (*Continued*)

Menu Entry	Description
Monitors	See "About the Monitors Page" on page 1247.
Tests	See "About the Tests Page" on page 1100.
Measurement Designer	See "About the Measurement Designer Page" on page 955.
Network (sub-menu)	
Topology	See "Network Visualization Options" on page 501.
Routing (sub-menu)	
Route Explorer	See "About the Route Explorer Page" on page 566.

Trust Menu

[Table 10 on page 26](#) displays the menu entries for the trust and compliance use case and links to relevant topics that you can refer to for more information.

Table 10: Trust Menu Entries

Menu Entry	Description
General (sub-menu)	
Trust Plans	See "About the Network Score Formula Page" on page 619.
Compliance (sub-menu)	
Compliance Scan	See "About the Compliance Page" on page 624.
Checklist	See "About the Compliance Checklist Page" on page 613.

Table 10: Trust Menu Entries *(Continued)*

Menu Entry	Description
Benchmarks	See "About the Compliance Benchmarks Page" on page 608.
Tailorings	See "About the Compliance Tailorings Page" on page 610.
Integrity (sub-menu)	
Hardware EOL	See "About the Hardware End of Life Page" on page 641.
Software EOL	See "About the Software Images Page" on page 363.
Vulnerabilities (sub-menu)	
Advisories	See "About the SIRT Advisories Page" on page 632.
PBNs	See "About the Proactive Bug Notifications Page" on page 635.

Orchestration Menu

[Table 11 on page 27](#) displays the menu entries for the trust and compliance use case and links to relevant topics that you can refer to for more information.

Table 11: Orchestration Menu Entries

Menu Entry	Description
Service (sub-menu)	
Instances	See "About the Service Instances Page" on page 829.
Service Catalog	See "About the Service Designs Page" on page 653.

Table 11: Orchestration Menu Entries (Continued)

Menu Entry	Description
Customers	See "About the Customer Inventory Page" on page 664.
Resource Instances	See "About the Resource Instances Page" on page 672.
Monitoring (Sub Menu)	
Workflows	See "About the Workflows Page" on page 880.
Service Orders	See "About the Service Orders Page" on page 874.

Inventory Menu

[Table 12 on page 28](#) displays the menu entries for device onboarding, device managements, Test Agents and sites, and links to relevant topics that you can refer to for more information.

Table 12: Inventory Menu Entries

Menu Entry	Description
Devices (sub-menu)	
Network Inventory	See "About the Inventory Page" on page 128.
Devices and Interface Profiles	See "About the Device and Interface Profiles Page" on page 205.
Configuration Templates	See "About the Configuration Templates Page " on page 376.
Configuration Backups	See "About the Configuration Backups Page" on page 372.
Software Images	See "About the Software Images Page" on page 363.

Table 12: Inventory Menu Entries *(Continued)*

Menu Entry	Description
Device Onboarding (sub-menu)	
Onboarding Dashboard	See "About the Put Devices into Service Page" on page 264.
Network Implementation Plan	See "About the Network Implementation Plan Page" on page 234.
Active Assurance	
Test Agents	See "About the Test Agents Page" on page 900.
Test Templates	See "About the Test Templates Page" on page 1288.
Common Resources	
Sites	See "About the Sites Page" on page 88.
Tags	See "About the Tags Page" on page 141.

Network Optimization

[Table 13 on page 29](#) displays the menu entry for the network optimization intent, and link to relevant topic that you can refer to for more information.

Table 13: Network Optimization Menu Entry

Menu Entry	Description
Network (sub-menu)	
Network Optimization Intent	See "About the Network Optimization Intent Page" on page 1298.

Planning

Table 14 on page 30 displays the menu entry for the planning, and link to relevant topic that you can refer to for more information.

Table 14: Planning Menu Entry

Menu Entry	Description
Networks (sub-menu)	
Offline Models	See "About the Working Models Page" on page 1336 .

RELATED DOCUMENTATION

[Routing Director Overview](#) | 3

Personas Overview

The management and operation of a network require different people to be involved at various stages of the process, and to perform tasks related to their area of expertise. This might mean that different departments handle different tasks, with handoffs between departments taking place. For example, one person might install a device, but a different person might then monitor the device onboarding process.

Routing Director is designed around a structured planning process that makes the life-cycle of the device and network efficient. By using structured planning, you can streamline the device onboarding and monitoring activities.

Routing Director uses personas to delineate the device life-cycle management (LCM) process. These personas provide a way for operators to map the different activities in the device LCM process to Routing Director.



NOTE: Personas are different from predefined *roles* that exist in the Routing Director GUI. Roles define which access permissions are available to users who are assigned to a role. However, a persona is simply a *logical* construct to make it easier to understand the structured planning approach for device LCM in Routing Director. For details about roles, see ["Predefined User Roles Overview" on page 97](#).

Table 15 on page 31 lists the different personas in Routing Director and the tasks that the persona performs.

Table 15: Personas in Routing Director

Persona	Description
Network Architect or Designer	A Network Architect typically performs the Day -2 activities in the device LCM process. These activities include: • Deciding the types of devices to be used in the network, and the configuration of the device types. • Identifying the types of interfaces to be used on different devices. • Determining what protocols need to run on the different types of devices. In addition, a Network Architect usually performs advanced troubleshooting tasks. In Routing Director, these tasks include creating resource pools, device profiles, interface profiles, and so on.
Network Planner (also known as Deployment Planner)	A Network Planner typically performs the Day -1 activities in the device LCM process. These activities include: • Defining the devices to be used and configuring the interfaces on the devices. • Defining how devices are connected and the topology to be used. In Routing Director, the Network Planner performs these tasks by creating a network implementation plan.

Table 15: Personas in Routing Director (*Continued*)

Persona	Description
Field Technician	A field technician typically performs the Day 0 activities in the device LCM process, These activities include: • Physical installation of the device. • Connecting the cables. • Inserting pluggables • Triggering the device onboarding. In Routing Director, the field technician uses a web-based GUI accessible on a handheld device or a laptop to perform the Day 0 activities.
NOC Engineer	A Network Operations Center (NOC) engineer oversees the Day 0 activities, and performs Day 1 activities and performs Day 2 activities. These activities include: • (Day 0 and Day 1) Monitoring the Day 0 activities of the field technician. Applying additional device configurations, and testing and certifying the device for production. • (Day 2 and beyond) Monitoring and troubleshooting devices, and so on.
IT or System Administrator	An IT or a System Administrator is involved only in the tasks related to the administration of Routing Director. This persona typically does <i>not</i> perform device LCM activities.
Life-cycle Manager	A life-cycle manager is an owner for a set a devices that plans software upgrades, handles device EOL, and works in tight collaboration with the Network Architect.

Table 15: Personas in Routing Director *(Continued)*

Persona	Description
SOC (Service Operation Center) Engineer	A SOC manager monitors one or more services in the network and take action to troubleshoot and remediate problems for specific customers.
Delivery Engineering	A team that executes or supervises the deliveries of new services to customers
Service Manager	Service Manager acts as counterpart to a specific customer (usually large) to report defects, SLAs, and so on.
First Line Support	The team that provides first level of support to customers and usually has restricted access to dedicated troubleshooting tools.
Second Line Support	The team that provides second level of support to customers and provides advanced support for devices.
Third Line Support	The team that provides third level of support to customers. This team works along with the Life-cycle Manager and the Network Architect.
Service Architect	Service Architect decides how devices should be configured to deliver a specific service to a customer. The service architect also does advanced troubleshooting for the customer's services.
Security Engineering	The team that owns all security aspects in the network.

For more information about the device LCM process, see ["Device Life-Cycle Management Overview" on page 168](#).

Retention Policies in Paragon Automation

SUMMARY

Use this topic to learn about the different retention policies used in Paragon Automation.

IN THIS SECTION

- [Observability | 34](#)
- [Service Orchestration | 34](#)
- [Active Assurance | 35](#)
- [Administration | 36](#)
- [Infrastructure | 36](#)

A retention policy is the period of time for which data is retained in the various databases before being purged. Retention policies in Paragon Automation vary by use case components and are dependent on the database and services used by the features. The following sections list the retention policies for the different use cases and components in Paragon Automation.

Observability

Table 16: Retention Policy for Observability Features

Feature	Retention Policy
KPIs and events	Stored for seven days.
Device logs	Stored for seven days.
Device alerts	Purged two hours after the alert is closed.
Device alarms	Stored in OpenSearch indefinitely or until the data is manually deleted.

Service Orchestration

[Table 17 on page 35](#) lists the retention policies used in service orchestration.

Table 17: Retention Policy for Service Orchestration

Feature	Retention Policy
Service instance	Stored until the service is deleted.
Service order history	Stored until the service is deleted.
Service order modifications	Stored in the audit log.
Device configuration modifications	Stored in the audit log
Workflow logs in Airflow	Stored in the PVC at /opt/airflow/mount/logs, with a default storage capacity of 10-GB. A DAG, <code>airflow_logs_cleanup.py</code>, is scheduled to run every 6 hours to trim the oldest logs when the PVC disk usage exceeds 80%. To retain logs for a longer period, increase the PVC size to allocate more storage space.

Active Assurance

Table 18 on page 35 lists information on the retention policy for metrics related to a stream.

Table 18: Retention Policy for Active Assurance

Metrics for a Stream	Retention Policy
Raw Data	3 days NOTE: Active Assurance does not store the raw data forever. As the data gets older, the data is stored at a lower resolution (concise version of a data). This process is called as data rollup. Data can be rolled up for a period of 1 minute, 5 minutes, 30 minutes, and 1 hour.
1 minute data rollup	7 days

Table 18: Retention Policy for Active Assurance (Continued)

Metrics for a Stream	Retention Policy
5 minutes data rollup	30 days
30 minutes data rollup	180 days

Administration

[Table 19 on page 36](#) lists the retention policies used by the administration features.

Table 19: Retention Policy for Administration Features

Data	Retention Policy
Audit logs	3 months
System logs	30 days
Device backups	2 months

Infrastructure

[Table 20 on page 37](#) lists the retention policies used by the Paragon Automation infrastructure.

Table 20: Retention Policy for the Infrastructure Databases

Data	Retention Policy
OpenSearch data	Indices with <code>*-_yyymmm</code> or <code>*__yyymmm</code> or <code>*_yyymmm-xxxxxxx</code> names are stored for the last three months. Indices with <code>*-_yyymmmdd</code> or <code>*__yyymmmdd</code> names are stored for the last 30 days. An automatic cleanup script runs every day at midnight to clean up indices that have exceeded their retention period. However, if the corresponding application cleans up the application data earlier, the application cleanup task will override the OpenSearch retention policy.

Access and Manage Routing Director Account

IN THIS CHAPTER

- [Access the Routing Director GUI | 38](#)
- [User Activation and Login | 39](#)
- [Reset Your Password | 42](#)

Access the Routing Director GUI

Routing Director provides you with multiple authentication methods to log in. The login workflow consists of the following tasks based on the authentication method that you choose.

To log in to Routing Director Web GUI:

1. Access the Routing Director Web GUI by entering the *web-ui-ip-address*.
2. Enter your credentials to log in to the Routing Director Web GUI. For more information on user login, see ["User Activation and Login" on page 39](#).
3. Create or select (join) an organization.

After you complete the login steps, you can view the Troubleshoot Devices page of an organization. You can secure your future login sessions to an organization by enabling two-factor authentication (2FA). If you enabled 2FA, you must verify your identity by using an authenticator application.

You can also configure single sign-on (SSO) that uses an identity provider (IdP) to authenticate and authorize users and to permit them to perform role-based tasks. For more information, see ["Single Sign-On" on page 66](#).

RELATED DOCUMENTATION

| [Authentication Methods Overview | 65](#)

User Activation and Login

IN THIS SECTION

- [Log into Routing Director for the First Time After Installation | 40](#)
- [Log in as a New User Without an Invite \(when SMTP is not Configured\) | 40](#)
- [Log in as a New User for the First Time Without an Invite \(when SMTP is Configured\) | 41](#)
- [Log in as a New User With an Invite \(when SMTP is Configured\) | 41](#)
- [Access a New Organization \(when SMTP is Configured\) | 42](#)

To access Routing Director Web GUI, you must have an account in Routing Director. Your account is activated after you log into the Routing Director Web GUI. You can then perform activities as defined by your role.



NOTE: If single sign-on (SSO) is enabled, you can access Routing Director through your identity provider (IdP) account. IdP authenticates your access privileges in Routing Director. For more information, see ["Single Sign-On" on page 66](#).

The system administrator who installs Routing Director is usually the first user to log in by using the e-mail address and password entered while installing Routing Director. Once logged in, the system administrator creates an organization and is assigned the Super User role in the organization. A superuser can add or invite users to the organization.

Depending on whether SMTP or local authentication is configured, Routing Director sends invitations and notifications to users. See ["Configure SMTP Settings in Paragon Shell" on page 101](#) for information about the behavior when either or both local authentication and SMTP are enabled or disabled.

If SMTP is configured in Routing Director, an invitation is sent over an e-mail when the superuser invites users to an organization. A user can click the link in the e-mail invitation and complete the login tasks. Your login procedure depends on whether you are an existing user with a Routing Director account or a new user without a Routing Director account.

If SMTP is not configured in Routing Director, the Super User adds users to an organization and shares the Web URL of Routing Director and credentials (e-mail ID and temporary password) to log in, with the users. Users can then access the Web URL and log in using the credentials. Once logged in to Routing Director, the user is prompted to create a new password.

After the user logs in and accesses an organization, the first page that Routing Director displays depends on your user role. If your role is Installer, the first GUI page you view is the Onboard a device page. For users with other roles, Routing Director displays the Troubleshoot Devices page.

The log in procedures for different roles are as follows:

Log into Routing Director for the First Time After Installation

The first user to log into Routing Director is typically the system administrator who installs Routing Director. On successful installation, the system administrator gets the URL to access the Routing Director Web GUI from Paragon Shell. The Welcome message that appears after logging in lists the IP address and e-mail ID to access the Routing Director Web GUI.

To log in as the first user:

1. Copy the Routing Director Web URL (for example, `https:// web-ui-ip-address`) from Paragon Shell, and enter the Web URL in a Web browser.

The Routing Director login page opens.



NOTE: Juniper Networks recommends that you use the latest version of Chrome, Firefox, or Safari browsers to access Routing Director.

2. Enter the e-mail address that appears in the Welcome message (the same e-mail address that you entered while installing Routing Director) after logging into Paragon Shell and click **Next**.
3. Enter the password that you configured during installation and click **Log in**.

The New Account page appears.

4. (Optional) Click **View Account** to check your user name and e-mail address.
5. Click **Create Organization**.

The Create Organization page appears.

Type a unique name for your organization and click **Create**.

An organization is created and you are logged into the organization as a superuser.

You can now add new users, create sites, and configure organization settings.

Log in as a New User Without an Invite (when SMTP is not Configured)

When SMTP is not configured in Routing Director, the superuser who creates your account in an organization manually informs you about the Web URL and the credentials (e-mail address and temporary password) to log into the Routing Director Web GUI.

To log in as a new user without an invite:

1. In a browser, enter the Routing Director Web URL shared by the Super User.

The Routing Director login page opens.

2. Enter the e-mail address (shared by the superuser) and click **Next**.
3. Enter the temporary password (shared by the superuser) and click **Log in**.

You are prompted to change your password.



NOTE: In case you forgot the password, please contact the superuser who shared the temporary password with you to reset and share the password.

4. Enter and re-enter the new password and click **Change Password**.

The password can contain up to 32 characters, including special characters, based on the password policy of the organization.

You have successfully logged in to Routing Director. The tasks you can perform in this organization depends on your user role. See "[Predefined User Roles Overview](#)" on page 97 for more information.

Log in as a New User for the First Time Without an Invite (when SMTP is Configured)

To access Routing Director Web GUI as a first-time user without an invite, you need the Web URL of Routing Director. To access the Web GUI, you must create an account in Routing Director. You will then receive a verification e-mail to validate your account.

To log in as the first user without an invite:

1. In a Web browser, enter the Routing Director Web URL (for example, `https:// web-ui-ip-address`).
2. Click **Create Account** on the login page.
3. Type your first name, last name, e-mail address, and password.

The password is case sensitive.

4. Click **Create Account**.

Routing Director sends a verification e-mail to activate your account.

5. Click **Validate me** in the e-mail body.

The New Account page appears.

6. (Optional) Click **View Account** to check your name and e-mail address.

7. Click **Create Organization**.

Type a unique name for your organization and click **Create**.

You are logged in as the superuser to the organization in Routing Director. See "[Predefined User Roles Overview](#)" on page 97 for information about tasks that you can perform in an organization as a superuser.

Log in as a New User With an Invite (when SMTP is Configured)

If SMTP is configured in Routing Director, you receive an e-mail invitation when the superuser creates an account for you to access an organization. You are also notified when your user account is modified.

To log in as a new user with an invite:

1. Click **Go to *organization-name*** in the e-mail body.

The Invite to Organization page opens in a Web browser.

2. Click **Register to accept**.

The My Account page appears.

3. Enter your first name, last name, e-mail address, and configure a password.

The password can contain up to 32 characters, including special characters, based on the password policy of the organization.

4. Click **Create Account**.

Routing Director sends a confirmation e-mail to activate your account.

5. In your confirmation e-mail, click **Validate me**.

The New Account page opens in a browser.

6. Click the organization for which you received the invite.

You can access the selected organization's GUI in Routing Director. The tasks you can perform in this organization depends on your user role. See "[Predefined User Roles Overview](#)" on page 97 for more information.

Access a New Organization (when SMTP is Configured)

To accept an invite to a new organization as an existing user already having an account in Routing Director, click the ***organization-name*** in the e-mail body. The Troubleshoot Devices (**Observability > Health > Troubleshoot Devices**) page of the organization opens in a browser.

The tasks you can perform in this organization depends on your user role. See "[Predefined User Roles Overview](#)" on page 97 for more information.

SEE ALSO

| [Organization and Sites Overview](#) | 57

Reset Your Password

IN THIS SECTION

- [Reset your Password \(When SMTP is Configured\)](#) | 43

● Reset Password (When SMTP is Not Configured) | 44

Routing Director allows you to reset your account password. You can reset your password from the My Accounts page after you access an organization, in the Routing Director Web GUI. See "[Manage Your Routing Director Account](#)" on page 125.



NOTE: If you had enabled two factor authentication for your account on the My Account page, it will be disabled when you reset your password. You must re-enable two factor authentication after logging into the GUI using your new password.

The workflow to reset your password when you are not logged into Routing Director varies depending on whether SMTP is configured or not in Routing Director.

Reset your Password (When SMTP is Configured)

IN THIS SECTION

● Reset your Password | 43

If SMTP is configured in Routing Director, all users have an option to reset the password from the login page.

Reset your Password

To reset the password from the login page:

1. On the Routing Director login page, type your e-mail address and click **Next**.
2. Click **Forgot Your Password?**

The Reset Password page appears.

3. Type your e-mail address and click **Send Reset Link**.

A message confirms that the link to reset password is sent to your e-mail address.

The Routing Director login page appears.

4. Click **Reset My Password** in the message body of the password recovery e-mail in your inbox.

The Set New Password page appears.

5. Type a new password in the Change Password text box and click **Change Password**.

A password must contain eight or more characters that are a combination of upper case and lower case letters, numbers 0-9, and special characters.

The Routing Director page appears.

6. Type your e-mail address and click **Next**.

The Routing Director login page appears.

7. Enter your new password and click **Log in**.

The Select an Organization page appears.

8. Join or create an organization.

You are logged into the Routing Director Web GUI.

Reset Password (When SMTP is Not Configured)

IN THIS SECTION

- [Reset the Password of a Superuser, in Paragon Shell, as a System Administrator | 44](#)
- [Reset the Password of a User from the Web GUI | 45](#)
- [Reset the Password as a User Logging in For the First Time | 45](#)

Depending on the type of user, different workflows exist to reset the password if SMTP is not configured in Routing Director.

The workflows are described in this section.

Reset the Password of a Superuser, in Paragon Shell, as a System Administrator

To reset the password of a superuser, in Paragon Shell, as a system administrator:

1. SSH as a root user to a node of the Routing Director cluster.

You are logged into Paragon Shell.

2. Execute the following command in the operational mode to reset the password.

```
request paragon super-user password reset user e-mail-address password temporary-new-password
```

where,

e-mail-address, is the e-mail address with which the superuser logs in to Routing Director.

temporary-new-password, is a temporary new password to be used for logging in to Routing Director.

The password is reset. You can now share the temporary new password with the superuser. The superuser can use the temporary password to access the Routing Director Web GUI. Upon logging in, the user will be prompted to reset the password.

Reset the Password of a User from the Web GUI

Superusers can reset the password of other users from the Users (**System Settings > Users**) page in the Routing Director Web GUI.

To reset the password of a user, from the Web GUI, as a superuser:

1. Click **Settings Menu > Users** on the banner.

The Users page appears.

2. Select the user whose password you want to reset and do one of the following:



NOTE: You can reset the password of users whose status is Created or Active.

- Click the **Edit User** (pencil) icon and then click **Reset Password** on the User: *User-Name* page.
- Click **More > Reset Password**.
- Right-click the user and click **Reset Password**.

The Reset Password page appears. You are presented with a randomly generated temporary password.

3. Click **Copy to Clipboard** to copy the masked password.

To view the masked password, click **Show Password**.

You must manually share the new password with the user for them to access the organization.

4. Click **OK**.

The user is listed on the Users page with the status as Created. The status changes to Active when the user successfully logs in to the application with the new password.

Reset the Password as a User Logging in For the First Time

The system administrator adds you to an organization and shares the Routing Director Web URL and credentials (e-mail address and temporary password) with you.

To reset the password as a user logging in for the first time:

1. Click the Routing Director Web URL.

Routing Director login page opens in your default browser.

2. Enter the e-mail address and click **Next**.

3. Enter the temporary password and click **Log in**.

You are prompted to change your password.

4. Enter and re-enter a new password and click **Change Password**.

You are logged into the Routing Director Web GUI.



NOTE: In case, you forgot your password and want to reset it, please contact the superuser who shared the temporary password. The superuser will reset and share a new temporary password.

SEE ALSO

| [Access the Routing Director GUI](#) | 38

CHAPTER 3

LLM Connector

IN THIS CHAPTER

- [LLM Connector Overview | 47](#)
- [Configure LLM Connector | 50](#)

LLM Connector Overview

SUMMARY

The LLM Connector tool, embedded in Routing Director, is an advanced AI-driven tool for leveraging your large language models (LLMs) (also known as bring your own LLM) to streamline network monitoring operations.

IN THIS SECTION

- [LLM Connector Actions | 48](#)
- [Benefits of LLM Connector | 49](#)



NOTE: The LLM Connector tool is a beta feature in this release.

LLM Connector facilitates the use of natural language to query network status and obtain troubleshooting information, without the need for traditional CLI commands. For example, if you encounter an issue with OSPF neighbors of a device, you query LLM Connector `What is the status of OSPF neighbors of device A?`. LLM Connector by using the knowledge that the LLM (defined and configured by the user) has obtained from technical documentation and by its ability to execute functions executes the `show ospf neighbor` command on the device to present you with the relevant information for troubleshooting.

When you initiate a query, LLM Connector connects with an LLM such as OpenAI GPT-4o, or Llama 3.1 that you configure to provide contextual assistance. The LLM is trained on extensive documentation to interpret your query and prompts LLM Connector to execute the necessary "read-only" network commands. The LLM also parses the result of the executed command and presents it in an easy-to-understand format.

The natural language interface supports a variety of query types, ranging from simple status checks to more complex troubleshooting commands. For example, you can ask, Show me the interface statistics or Show insights on input traffic rate on et-0/0/0 on device A.

LLM Connector can help you with:

- Retrieving device information.
- Executing Junos OS operational commands.
- Getting Insights based on the telemetry collected from the device.
- Retrieving a list of all VPNs in your network and their details, metrics, and health information.
- Fetching information about customers and service instances associated with customers.
- Plotting graphs related to KPIs. For example, plot graphs for CPU utilization, fan RPM, and device temperature.

To use the LLM Connector tool, you need to set up the LLM and provide parameters such as API keys and model details. For information about configuring LLM, see ["Configure LLM Connector" on page 50](#).

You can rate and provide feedback regarding the output response. Routing Director stores your feedback for up to an year.

You can access LLM Connector by clicking the LLM Connector icon displayed on the right-bottom corner of the UI. However, you can move the icon around and anchor it on any of the four corners on the GUI.

LLM Connector provides a chat window where you can enter your queries. LLM Connector retains your queries in a session so that you can continue with the queries after a break, if needed. A session is active for an hour. If your session is idle for more than an hour, LLM Connector creates a new session for your next set of queries. If you want to go back to any of your older queries, click the **Chat History** (clock) icon (present on right-corner of the LLM Connector window). The chat history lists all your previous conversations. Click on any conversation to continue the conversation.

LLM Connector Actions

LLM Connector provides the following action menu items on the LLM Connector window:

- **Chat History**—Click the **Chat History** (clock) icon on the top of the LLM Connector window to view chat history. Click the Chat History icon again to close the chat history.

The chat history is listed, grouped by the time of chat as Today, Yesterday, Last Week, and so on. Chat History is displayed on the left side of the LLM Connector window when the window is maximized.

To open a new conversation, click the **New Conversation** icon (+).

- **Search**—Enter a search text in the **Search** field to search for conversations and chats. LLM Connector matches the search text that you enter with the your query text in the chat history and fetches the corresponding conversation.

The Search field is present above the Chat History list and displayed when you view the chat history.

- **Settings**—Configure the following settings for the LLM Connector:

- **Model**—Select the model for use by LLM Connector.

Available options: OpenAI or Ollama

- **Temperature**—Set the LLM temperature. LLM Temperature helps the LLM to produce more coherent and consistent text and avoid irrelevant responses.

Range: 0.1 (default) to 1

- **Top P**—Set the Top P of the LLM connector. LLM Top P helps the LLM to avoid providing inconsistent outputs.

- Range: 0.1 (default) to 1

- **Maximize**—Click to maximize the LLM Connector window. When the LLM Connector window is maximized, chat history is listed on the left side of the window.

- **Close**—Click to close the LLM Connector window.

Benefits of LLM Connector

- **Reduces Workload on NOC Teams:** By automating repetitive troubleshooting tasks, LLM Connector significantly lowers the operational burden on Network Operations Center (NOC) teams, freeing them to focus on more complex issues.
- **Accelerates Problem Resolution:** LLM Connector provides rapid, contextually relevant solutions by executing necessary commands and parsing results, which enables you to solve problems quickly.
- **Simplifies Training Requirements:** LLM Connector minimizes the learning curve to work with Routing Director for new operators by allowing them to interact with the Routing Director using natural language queries, reducing the time and effort required to become proficient.
- **Enhances Data Privacy and Security:** By supporting Cloud, Private Cloud, and On-Premises deployments, LLM Connector offers flexible options that cater to various data privacy and security needs, ensuring secure operations.
- **Improves System Usability:** The integration of function calling capabilities with a range of LLM models enables precise and efficient interactions with external systems, enhancing the overall usability of Routing Director as a troubleshooting tool.

Configure LLM Connector

SUMMARY

Use this topic to configure LLM connector in Routing Director.

LLM connector connects large language models (LLMs) to the Routing Director infrastructure. This connection allows a user to query their network for troubleshooting and monitoring through Routing Director by using conversational English, bypassing the need for traditional CLI commands. For example, What is the status of the OSPF neighbors of *device-a*?

To be able to query your network through Routing Director, you should configure the LLM connector in Routing Director.

To configure LLM Connector:

- 1. Click **Settings Menu > System Settings** on the banner.
The Organization Settings page appears.
- 2. On the Configure LLM Connector tile, click + (Add) icon.
The LLM Connector Configuration page appears.
- 3. Enter values referring to table [Table 21 on page 50](#).
Fields marked * are mandatory.
- 4. Click **Create**.
The values are listed in the Configure LLM Connector tile.

Table 21: Fields to configure LLM connector

Field	Description
Name*	Enter a name for the LLM model. The name need not be unique within an organization.
Provider*	Select the provider of the LLM. For example, Azure Open AI, Open AI, and Ollama. Open AI is the recommended LLM provider in this release.

Table 21: Fields to configure LLM connector (*Continued*)

Field	Description
Model*	Select the LLM to be used. For example, GPT-3.5, GPT-4, GPT-4o, or Llama3.1. GPT-4 and GPT4o are the recommended models to be used in this release.
API Key*	Enter an API key for the LLM. The LLM uses the API key to authenticate requests from LLM Connector.
Base URL*	Enter the base URL value of the Azure Open AI model and Ollama. You obtain the base URL when you deploy your LLM. This is a mandatory field if you use Azure or Ollama as your LLM provider.
API Version*	Enter the API version of the AI model. You obtain the base URL when you deploy your LLM. This field is mandatory for Azure.
Active	Click to enable (default) or disable the LLM. Only one LLM configuration can be active at a time.

2

PART

Administration

- Introduction | 53
 - Organization Management | 57
 - Site Management | 88
 - User Management | 93
 - Inventory Management | 128
 - Audit Logs | 137
 - Tags | 141
 - Attribute-Based Access Control | 145
-

Introduction

IN THIS CHAPTER

- [Administration Overview | 53](#)
- [Administration Workflow | 55](#)

Administration Overview

IN THIS SECTION

- [Manage Organizations | 53](#)
- [Manage Sites | 54](#)
- [Manage Users | 54](#)
- [Monitor Audit Logs | 55](#)

Routing Director provides an easy-to-use user and organization management system. An administrator with the Super User role can manage organizations, sites, and the users in the organization. The user who creates the organization is assigned the Super User role (superuser) in the organization, by default. After the organization is created, the superuser can configure organization settings, add sites, and assign predefined roles to users according to the tasks the users need to perform in the organization. This topic provides an overview of the tasks a superuser performs in an organization.

Manage Organizations

After creating an organization, the superuser can configure the following features from the Settings page to efficiently manage the organization:

- Authentication methods to manage access to the organization

- Identity providers (IdP) to enable single sign-on (SSO)
- Map IdP user groups to predefined roles in Routing Director
- Session policy to time out sessions following a period of inactivity
- API tokens to enable users to retrieve information through REST APIs
- Password policy to secure users' access to Routing Director
- Webhooks to view alerts and events notifications in real-time
- RADIUS to authenticate devices during onboarding
- LLM configuration settings to facilitate users to LLM connector chatbot.
- SNMP manager to receive SNMP traps from devices and configure external applications to forward the received SNMP traps.
- Network optimization settings to automatically reroute LSPs to set database pruning frequency and postgres retention period.

For more information, see ["Organization and Sites Overview" on page 57](#).

Manage Sites

After you create an organization, you must create sites. Sites are physical locations within the organization. Sites house the devices in a network, such as routers, switches, and firewalls. After sites are created, a superuser can assign devices to those sites. The Sites page provides information about sites, their location and timezone, and the site group to which the sites belong. A superuser can add a site, edit the site information or delete one or more sites that are not in use.

For more information, see ["About the Sites Page" on page 88](#).

Manage Users

To perform the various tasks in an organization, the superuser must add users to various predefined roles according to the tasks the users must perform in the organization. Adding a user to the organization is as easy as entering the user's name, e-mail address, and assigning a predefined role in the organization. Based on the tasks that a user must perform, a superuser can assign roles, such as Super User, Network Admin, Observer, and Installer, providing role-based access to resources. A superuser can add, modify, and delete users. For more information, see ["About the Users Page" on page 93](#).

Monitor Audit Logs

An audit log is a record of a sequence of activities initiated by a user or by a process in a workflow that the user has initiated. Routing Director stores audit logs for 30 days. Audit logs are useful in tracking and maintaining the history of users' activities in Routing Director. For more information, see ["About the Audit Logs Page" on page 138](#).

RELATED DOCUMENTATION

[Manage Organization Settings | 60](#)

[LLM Connector Overview | 47](#)

[Network Optimization Overview | 1294](#)

Administration Workflow

After Routing Director is installed, tasks related to the administration of Routing Director can be performed. The first user who accesses Routing Director is assigned the Super User role by default.

After logging in, the superuser must create an organization, which consists of users, devices, and geographical sites in the network. Next, the superuser must perform administration tasks. ["Figure 1" on page 55](#) shows the sequence of tasks that the superuser performs, starting with user account creation.

Figure 8: Administration Workflow



The tasks that a superuser must perform are as follows:

1. Log in to Routing Director and create an account.

See ["User Activation and Login" on page 39](#).

2. Create an organization.

See ["Add an Organization" on page 58](#).

3. Configure organization settings—You must configure the following for your organization:

- Password policy
- Single sign-on (SSO) if you want to authenticate and to authorize users by using an Identity Provider (IdP)

You can optionally configure other organization settings such as session and inactivity timeouts, API tokens, and so on.

See ["Manage Organization Settings" on page 60](#).

4. Invite users to the organization—You can invite users in either of the following ways:

- By adding a user to an organization and assigning a role to the user. The tasks that a user performs depends on the assigned role. See ["Invite Users" on page 109](#) to send invites and ["Manage Users and Invites" on page 112](#) to manage users and invites in an organization.
- By configuring a third-party IdP that authenticates and authorizes users based on the role mapped to each user group. See ["Manage Identity Providers" on page 68](#).

5. Create one or more sites—A site represents a geographical location with one or more devices in your network. However, a device can be associated with only one site. See ["Add Sites" on page 90](#).

After you perform the initial administration related tasks, you can explore other tasks such as viewing inventory (**Inventory > Network Inventory**) and monitoring audit logs (**Settings Menu > Audit Logs**). See ["About the Inventory Page" on page 128](#) and ["About the Audit Logs Page" on page 138](#).

RELATED DOCUMENTATION

| [Audit Logs Overview](#) | 137

Organization Management

IN THIS CHAPTER


- [Organization and Sites Overview | 57](#)
- [Add an Organization | 58](#)
- [Delete an Organization | 59](#)
- [Manage Organization Settings | 60](#)
- [Authentication Methods Overview | 65](#)
- [Manage Identity Providers | 68](#)
- [Manage API Tokens | 72](#)
- [Manage Roles | 75](#)
- [Configure SNMP | 77](#)
- [Configure Webhooks to Receive Event Notifications | 81](#)
- [Manage RADIUS Server Configurations | 84](#)

Organization and Sites Overview

An organization in Routing Director represents a customer. An organization can have multiple sites representing the locations where routers, switches, and firewalls are installed. While a site can have more than one device, a device can be associated with only one site. In Routing Director, you must assign a device to a site to be able to apply the device life-cycle management (LCM) functions on the device. You can add only one organization in this release.

You can group sites based on regions, functions, or other parameters for efficient management of the devices. [Figure 9 on page 58](#) represents the relation between an organization, sites, and site groups in Routing Director. In [Figure 9 on page 58](#), an organization has seven sites and three sites groups (Site Group 1, Site Group 2, and Site Group 3). Site 3 and Site 4 are a part of Site Group 1 and Site Group 3 while Site 7 is part of Site Group 2 and Site Group 3.

Figure 9: Organization, Sites, and Site Groups



RELATED DOCUMENTATION

[Add Sites](#) | 90

Add an Organization

An organization represents a customer. You can add only one organization in this release.

You can add an organization from:

- The New Account page when you log in to Routing Director as a superuser.
- The organization list (next to the Help icon) on the top right-corner of the Routing Director GUI.

To add an organization:

1. Click **Create Organization** on the New Account page or in the Organization drop-down list at the top-right corner.

The Create Organization page appears.

2. In the **Organization Name** field, enter a name for the organization.
3. Click **OK**.
The organization appears in the list of organizations and on the Select an organization page.
4. Click the organization to access the organization.

You are the superuser for an organization that you create. After you create an organization, you can configure the organization settings and invite users to access the organization. For more information, see ["Manage Organization Settings" on page 60](#) and ["Invite Users" on page 109](#) respectively.

RELATED DOCUMENTATION

[Administration Workflow | 55](#)

[Add Sites | 90](#)

Delete an Organization

You can delete an organization that you no longer manage or if you want to decommission the organization. You must be a user with the Super User role to delete an organization.



CAUTION: You cannot restore an organization after you delete it.

To delete an organization:

1. Log in to Routing Director and click the organization that you want to delete.
2. Click **Settings Menu > System Settings** on the banner.
The Organization Settings page is displayed.
3. Click **Delete Organization**.
The Delete Organization page appears.
4. As a confirmation for deleting the organization, enter the name of the organization in the **Organization Name** field.
5. Click **Delete Organization**.
The organization is deleted and the Routing Director login page appears.

RELATED DOCUMENTATION

[Organization and Sites Overview | 57](#)

[Administration Workflow | 55](#)

Manage Organization Settings

A superuser can configure the organization settings and do the following tasks:

- View organization name and organization ID and modify the organization name.
- Add, modify, and delete identity providers (IdP).
- Add, modify, and delete custom roles.
- Enable or disable the password policy for the organization and modify the password policy when the password policy is enabled.
- Modify the session timeout policy for the organization.
- Generate, edit, and delete API tokens for various roles in the organization.
- Configure webhooks for the organization.
- Configure device authentication.
- Configure SNMP manager in Routing Director and the external applications to which the SNMP traps can be sent.
- Configure LLM connector settings.
- Configure Network Optimization settings.

To configure and to manage organization settings:

1. Click **Settings Menu > System Settings** on the banner.

The Organization Settings page appears.

2. Configure or modify the organization settings as needed. Refer to [Table 22 on page 61](#).

Fields marked * are mandatory.

3. Click **Save** to save the settings.

Verify that the settings are saved and close the Organization Settings page.

[Table 22 on page 61](#) describes the parameters on the Organization Settings page.

Table 22: Organization Settings Parameters

Field	Description
Organization Name*	Name of the organization. You can edit the organization name here.
Organization ID	The ID for the organization. The value is auto-generated. This is a read-only field.
Single Sign On (SSO)	
Identity Providers	View identity providers configured in the organization. Add, edit, or delete the identity providers; see "Manage Identity Providers" on page 68 .
Roles	View roles configured for SSO. Add, edit, or delete the roles; see "Manage Roles" on page 75 .
Password Policy	Enable or disable (default) password policy. If you enable the password policy, configure the password policy parameters; see Table 23 on page 62 .
Session Policy	Configure the time, in minutes, after which the session with Routing Director should timeout; see Table 24 on page 63 .
API Tokens	Generate and view API tokens to authenticate users when they retrieve data by using REST APIs; see "Manage API Tokens" on page 72 .
Webhooks	Webhooks enable you to get notifications when the events that you have subscribed for occur. Internet connectivity is required for Routing Director to connect to third-party applications, such as Slack. Click to enable or disable (default) webhooks. If you enable webhooks, you must select the type of events for which you want to receive notifications; see Table 25 on page 63 .
Device Authentication	Configure the type of authentication Routing Director must use to onboard a device—Local (default) or RADIUS. See "Manage RADIUS Server Configurations" on page 84 to configure RADIUS authentication.

Table 22: Organization Settings Parameters *(Continued)*

Field	Description
SNMP Forwarding Endpoints	Configure Routing Director to forward SNMP traps received from Juniper devices to external applications. See "Configure SNMP" on page 77.
SNMP Configuration	Configure the SNMP manager to receive SNMP traps from Juniper devices and forward them to external applications. See "Configure SNMP" on page 77.
Configure LLM Connector	Configure the LLM Connector to facilitate the use of natural language to query network status and obtain troubleshooting information. See "Configure LLM Connector" on page 50. NOTE: LLM Connector is a beta feature in this release.
Network Optimization Settings	Configure pathfinder settings to automatically reroute LSPs and to set database pruning frequency and Postgres retention period. See Table 26 on page 64 and "Reroute LSPs" on page 550.

Table 23: Parameters to Configure Password Policy

Field	Description
Required minimum password length*	Enter the minimum number of characters that should be present in the password of a user's account. Default is 8 characters. Range: 8 to 32
Require special characters*	Click to enable (default) or disable the use of special characters in the password.
Require 2-Factor Authentication*	Click to enable or disable (default) two-factor authentication for users accessing the organization. If you enable two-factor authentication, a code is sent to an authenticator app. The code should be entered in addition to the password to access the organization.

Table 24: Parameters to Configure Session Policy

Field	Description
Session Timeout (minutes)*	Enter the number of minutes after which the session should timeout. Default is 20,160 minutes.
Inactivity Timeout (minutes)*	Enter the number of inactive minutes after which the session should timeout. Default is 0, indicating that the session does not time out because of inactivity. Range: 0 to 480 minutes

Table 25: Parameters to Configure Webhooks

Field	Description
Name*	Enter the name of the server or application to which notifications for subscribed events are to be sent.
URL*	Enter the URL of the server or application where the notifications in the form of HTTP POST requests are to be sent when a subscribed event occurs. You must configure webhooks to enable Routing Director to send notifications to third-party applications, such as Slack, when events you have subscribed to are triggered on the managed devices. To receive webhook notifications in a format that is compatible with Slack, you need to configure an intermediary that can interact with the sending and receiving applications, in this case, Routing Director and Slack. The recommended intermediary platform is Make. For more information, see "Configure Webhooks to Receive Event Notifications" on page 81.
Secret	Enter the secret to validate that the notifications received are from valid hosts.
Webhook Header	
Header Key*	Enter a unique key that the webhook endpoint can use to authenticate the event notifications.
Header Value*	Enter a unique value for the key.
Streaming API	

Table 25: Parameters to Configure Webhooks (*Continued*)

Field	Description
Alerts	Click to enable or disable (default) receiving notifications when subscribed alerts are generated on the managed devices.
Audits	Click to enable or disable (default) receiving notifications when an organization is accessed or any setting in the organization is changed.
Device Status	Click to enable or disable (default) receiving notifications when the device status changes due to events such as a link going up or down, or the device getting disconnected from Routing Director, and so on.
Device Alarms	Click to enable or disable (default) receiving notifications when subscribed alarms are generated on the managed devices.

Table 26: Parameters to Configure Network Optimization Settings

Field	Description
ECMP Placement Method*	Select a placement method for equal-cost multipath (ECMP) routing. • Random—Randomly selects one of the ECMP paths • Least filled—Selects the path with the maximum available bandwidth • Most filled—Selects the path with the minimum available bandwidth Based on the placement method you opt, Routing Director evaluates link metrics, distributes traffic, and improves bandwidth utilization on multiple links to the same destination.
Path Optimization Interval*	Specify the interval after which Routing Director must automatically trigger path optimization. The path optimization is based on the current network and not on the most recent Path Analysis report. To disable automatic path optimization, enter 0 second. Range: 0 through 9 minutes (m), hours (h) or days (d).

Table 26: Parameters to Configure Network Optimization Settings (*Continued*)

Field	Description
Link Utilization Threshold*	Specify the threshold value (in percentage) for link utilization. When traffic on a link exceeds this value, Routing Director reroutes label-switched path (LSPs). If you do not specify a threshold, the LSP is not rerouted. If you specify 0, the links are blocked. Range: 0 through 100 percentage.
Threshold Rerouting Interval	Specify the minimum interval (in minutes) after which Routing Director reacts to any traffic or delay violations. If you don't specify an interval, the label-switched path (LSP) is not rerouted in case of a violation. Range: 1 through 300 minutes
PruneDB Cleanup Frequency*	Specify how often (in days) the database pruning (clean up) must happen. Range: 1 through 180 days
Postgres Retention*	Specify a period (in days) for which Postgres data must be retained before being purged. Range: 1 through 180 days.

RELATED DOCUMENTATION

[Organization and Sites Overview | 57](#)

[Administration Workflow | 55](#)

Authentication Methods Overview

IN THIS SECTION

- [Username and Password Authentication | 66](#)
- [Single Sign-On | 66](#)

Routing Director can authenticate users by using different authentication methods.

You can use one of the authentication methods that are listed in this topic to log in to the Routing Director Web GUI.

Username and Password Authentication

Users can create a Routing Director account to access the Routing Director Web GUI.

In the username and password authentication, a user enters the log in credentials in the Login page of Routing Director. Routing Director authenticates the identity of users by verifying the login credentials (username and password) entered by the users against the information stored in Routing Director database. This ensures that only users with valid credentials access Routing Director. For more information, see ["User Activation and Login" on page 39](#).

Single Sign-On

Routing Director can authenticate users by using single sign-on (SSO). SSO simplifies password management for users and administrators through centralized authentication by an identity provider (IdP) (for example, OpenLDAP, Microsoft Active Directory, and JumpCloud).

A superuser can configure IdP in the Organization Settings page and map default roles in Routing Director to the IdP user groups. Routing Director supports Secure Assertion Markup Language (SAML 2.0) for SSO authentication using IdPs and Lightweight Directory Access Protocol (LDAP). The IdP asserts a user's identity and allows the user to access the Web GUI based on the user's role.

To configure SSO in Routing Director:

1. Add the IdP to Routing Director; see ["Manage Identity Providers" on page 68](#).
2. Map users logging in by using the IdP account credentials to the predefined roles in Routing Director; see ["Manage Roles" on page 75](#).

After IdP is configured, superuser shares the SSO URL with the users.

An SSO login can be of two types: IdP-initiated and Service Provider (SP- initiated).

When using IdP-initiated login, a user must log in to the IdP and click or select the application they want to access. The IdP sends the authentication request as an XML-based SAML assertion to Routing Director. The user is logged in to Routing Director if Routing Director accepts the SAML assertion. Otherwise, the user's log in attempt to Routing Director fails.

When using SP-initiated login, the user accesses the Login page of Routing Director to enter their credentials. Routing Director redirects the authentication request to the IdP, to which the IdP responds with a SAML assertion. Routing Director accepts the SAML assertion from the IdP and logs in the user.



NOTE: If SSO is enabled, ensure that you use IdP-initiated login to log in to Routing Director. Using SP-initiated login will result in an error.

To sign in using SSO for the first time:

1. User must enter the SSO URL in a browser.

The login screen of the IdP appears.

IdP server authenticates the user based on the sign-in method configured. For example, an approval notification is sent to the user's registered device.

2. After the IdP server successfully authenticates the user, the user is logged in to the Routing Director Web GUI. Routing Director enforces access control on the user based on the role that the Routing Director superuser previously assigned for the IdP user group to which the user belongs.

Once a user is successfully authenticated, the user can avoid the process of repeated logins to access the Routing Director Web GUI. The user remains signed in until the authentication session expires.

RADIUS Authentication and Authorization



NOTE: In this release, RADIUS authentication and authorization support is limited to configuring RADIUS on new devices onboarded to Routing Director and authenticating and authorizing the device during onboarding.

By using RADIUS, you can onboard devices by using credentials of users configured in a RADIUS server.

To use RADIUS, you (superuser) must:

1. Install at least one RADIUS server in your network.
2. Enable RADIUS authentication in Routing Director.
3. Configure Routing Director to use the RADIUS server.

By default, local authentication is used when adopting a device. See ["Adopt a Juniper Device" on page 189](#). The adoption commands create a local user on the device, using generated credentials. On first SSH connection with the device, Routing Director updates the local password on the device with a newly generated credential.

When you enable RADIUS, Routing Director provides a set of outbound SSH commands to configure RADIUS on the device. After the device connects to Routing Director during onboarding, the device connects with the RADIUS server to authenticate the credentials provided by Routing Director.

To enable and configure a RADIUS server in Routing Director, see ["Manage RADIUS Server Configurations" on page 84](#).

RELATED DOCUMENTATION

[Administration Workflow](#) | 55

[User Activation and Login](#) | 39

Manage Identity Providers

IN THIS SECTION

- [Add an Identity Provider](#) | 71
- [Edit an Identity Provider](#) | 71
- [Delete an Identity Provider](#) | 71

Routing Director provides the Security Assertion Markup Language (SAML) and Lightweight Directory Access Protocol (LDAP) options to add identity providers (IdPs) for authentication of users logging into Routing Director.



NOTE:

- Active Directory is the only directory service supported in this release.
- You must map the user groups defined in the LDAP server to the roles in Paragon Automation along with configuring LDAP in Routing Director. For information about mapping user groups to roles, see ["Manage Roles" on page 75](#).

[Table 27 on page 69](#) lists the parameters to add identity providers to an organization.

Table 27: Parameters to Add Identity Providers

Field	Description
Name*	Enter a name for the IdP.
Type*	Select the type of IdP. The available options are: • SAML (default)—Use this option to allow users to log in by using third-party credentials; for example, credentials to log in to Google. See "SAML Options" on page 69 to configure SAML in Routing Director. • LDAP—Use this option if you have an LDAP directory storing the user authentication and authorization information. See "LDAP Options" on page 70 to configure LDAP in Routing Director.
SAML Options	
Issuer*	Enter the unique URL that identifies your SAML IdP. For example, Google or Microsoft. NOTE: Ensure that Routing Director is registered with the IdP so that you get the values to input for Issuer.
Name ID Format*	Select a unique ID for the user. The options are e-mail and unspecified. If you select e-mail, the IdP uses your e-mail address to authenticate you. If you select unspecified, the IdP generates a unique ID to authenticate you.
Signing Algorithm*	Select a signing algorithm from the following: • SHA1 • SHA256 (default) • SHA384 • SHA512
Certificate*	Enter the certificate issued by the SAML IdP. NOTE: Ensure that Routing Director is registered with the IdP so that you get the values to input for Certificate.

Table 27: Parameters to Add Identity Providers (*Continued*)

Field	Description
SSO URL*	Enter the URL to redirect the users to the SAML identity provider for authentication.
Custom Logout URL	Enter the URL to redirect the users after logging out.
ACS URL	The URL that the IdP should redirect an authenticated user to after signing in. The value is auto-generated and not editable.
Single Logout URL	The URL that the IdP should redirect when a user logs out of an authentication session. The value is auto-generated and not editable.
LDAP Options	
Server Host*	Enter the hostname of the LDAP server.
Server Port*	Enter a port number for the LDAP service. Range - 1 through 65,000
Base DN*	Enter the base distinguished name (DN) within the information tree in the LDAP server. The Base DN is the root tree for LDAP searches. For example: DomainComponent (dc)=company,dc=com
Manager DN	Enter the LDAP account (in full DN) for querying a user record for password verification and group association. Use this option when the LDAP server is configured to query with a password.
Manager Password	Enter a password for the user specified in the Manager DN field.
Verify SSL Cert	Click the check box to indicate whether the certificate of the LDAP server is to be validated.

Table 27: Parameters to Add Identity Providers (*Continued*)

Field	Description
CA Certificates	Enter the Privacy-Enhanced Mail (PEM)-coded certificates from the certificate authority. Enter each certificate in a new line.

Add an Identity Provider

To add an IdP:

1. Click **Settings Menu > System Settings** on the banner.
The Organization Settings page appears.
2. Click the **Create IDP (+)** icon above the Identity Providers table.
The Create Identity Provider page appears.
3. Configure the IdP by using the guidelines in [Table 27 on page 69](#).
Fields marked * are mandatory.
4. Click **Create**.
The IdP is created and listed in the Identity Providers table.

Edit an Identity Provider

To edit an IdP:

1. Click **Settings Menu > System Settings** on the banner.
The Organization Settings page appears.
2. Click the IdP you want to edit in the Identity Providers table.
The Edit Identity Provider page appears.
3. Edit the IdP by using the guidelines in [Table 27 on page 69](#).



NOTE: You cannot edit IdP type, ACS URL, and Single Logout URL.

4. Click **Save**.
You are returned to the Organization Settings page, where you can view the changes in Identity Providers table.

Delete an Identity Provider

To delete an IdP:

1. Click **Settings Menu > System Settings** on the banner.
The Organization Settings page appears.

- 2. Click the IdP that you want to delete.
The Edit Identity Provider page appears.
- 3. Click **Delete**.
You are returned to the Organization Settings page, where you can view that the IdP is removed from the Identity Providers table.

SEE ALSO

Administration Workflow	55
User Activation and Login	39

Manage API Tokens

IN THIS SECTION

- [Add an API Token](#) | [73](#)
- [Edit an API Token](#) | [74](#)
- [Delete an API Token](#) | [74](#)

API tokens authenticate users when they try to retrieve information from Routing Director by using REST APIs. By using API tokens, users can avoid authentication for each request that they make. An API token provides visibility into the resources accessed by a user, enabling you to have better control over access to resources.

[Table 28 on page 72](#) lists the parameters for configuring API tokens.

Table 28: Parameters to Configure API Tokens

Field	Description
Name*	Name of the API token.

Table 28: Parameters to Configure API Tokens *(Continued)*

Field	Description
Role*	Role to which the API token is applicable: • Super User • Network Admin You must assign at least one access control profile to a Network Admin role to grant access to resources. See "Access Control Profile Overview" on page 145. • Observer • Installer
Key	The key auto-generated to identify the application the user is using to access the resources.
Tags	Tags assigned to the API token. If there is more than one tag assigned to the API token, an integer is present next to the displayed tag. For example, +2 in the Tags column indicates that there are two more tags in addition to what is currently listed. You can hover over the integer to view the additional tags assigned to the API token. To assign a tag, edit the API token. See "Edit an API Token" on page 74.
Access Control Profile	Assign one or more access control profiles to grant access to resources (devices, service instances, and perform tagging operations). If an access profile does not exist, click the +Add Access Control Profile link to create a new profile and assign it to the Network Admin user. See "Add Access Control Profiles" on page 154.

Add an API Token

To add an API token for a role:

1. Click **Settings Menu > System Settings** on the banner.
The Organization Settings page appears.
2. Click the **Create Token (+)** icon.
The Create API Tokens page appears.
3. Enter values by following the guidelines in [Table 28 on page 72](#).

Fields marked * are mandatory.

4. Click **Generate**.

The API token is populated in the **Key** field.

5. Click **Close** to return to the Organization Settings page.

Edit an API Token

To edit an API token:

1. Click **Settings Menu > System Settings** on the banner.

The Organization Settings page appears.

2. Click the API token that you want to edit.

The Edit API Token page appears.

3. Add a pre-defined tag and edit the name, role, and assigned tags. If a tag does not exist, click the **Create New Tag** link to create a new tag and then assign the tag to the API token.

4. Click **Save**.

You are returned to the Organization Settings page, where you can verify the changes in the API Tokens table.

Delete an API Token

To delete an API token:



NOTE: Users using API tokens to access Routing Director resources cannot access the resources after the API token is deleted.

1. Click **Settings Menu > System Settings** on the banner.

The Organization Settings page appears.

2. Click the API token that you want to delete.

The Edit API token page appears.

3. Click **Delete**.

You are returned to the Organization Settings page, where you can verify that the API token is not listed in the API Tokens table.

SEE ALSO

[Access Control Profile Overview | 145](#)

[Administration Workflow | 55](#)

Manage Roles

IN THIS SECTION

- [Add a Role | 76](#)
- [Edit a Role | 77](#)
- [Delete a Role | 77](#)

A user with the Super User role can map predefined roles in Routing Director to the IdP user groups. For example, you can map an IdP group to the Network Admin role so that the users that are part of the group have the access privileges of the Network Admin user in Routing Director.



NOTE: If a user's role is defined both as part of an IdP and as a local user in Routing Director, then the access privileges assigned as a local user overrides access privileges defined for the IdP.

[Table 29 on page 75](#) lists the parameters to add custom roles to an organization.

Table 29: Parameters to Add Roles

Field	Description
Name*	Enter the name of the IdP user group.

Table 29: Parameters to Add Roles *(Continued)*

Field	Description
Role*	Select an access level for the user group: • Super User • Network Admin You must assign at least one access control profile to a Network Admin role to grant access to resources. See "Access Control Profile Overview" on page 145. • Observer (default) • Installer See "Predefined User Roles Overview" on page 97 for details on privileges of each role.
Access Control Profile*	Assign one or more access control profiles to grant access to resources (devices, service instances, and perform tagging operations). If an access profile does not exist, click the +Add Access Control Profile link to create a new profile and assign it to the Network Admin user. See "Add Access Control Profiles" on page 154.

Add a Role

A superuser can map an IdP user group to a pre-defined role in Routing Director.

To map an IdP user group to a pre-defined role:

1. Click **Settings Menu > System Settings** on the banner.
The Organization Settings page appears.
2. Click the **Create Role (+)** icon.
The Create Role page appears.
3. Configure the role by following the guidelines in [Table 29 on page 75](#).
Fields marked * are mandatory.
4. Click **Create**.
The role is listed in the Roles table.

Edit a Role

To modify the mapping of an IdP user group to a pre-defined role:

1. Click **Settings Menu > System Settings** on the banner.
The Organization Settings page appears.
2. Click the role that you want to edit.
The Edit Role page appears.
3. Edit the name and role by following the guidelines in [Table 29 on page 75](#).
4. Click **Save**.
You are returned to the Organization Settings page, where you can verify the changes in the Roles table.

Delete a Role

To delete the mapping of an IdP user group to a pre-defined role:

1. Click **Settings Menu > System Settings** on the banner.
The Organization Settings page appears.
2. Click the role that you want to delete.
The Edit Role page appears.
3. Click **Delete**.
You are returned to the Organization Settings page, where you can verify that the custom role is not listed in the Roles table.

SEE ALSO

[Add Users to an Organization Overview | 107](#)

[Predefined User Roles Overview | 97](#)

Configure SNMP

SUMMARY

Use this topic to configure Routing Director to send SNMP traps received from Juniper devices.

Routing Director receives and stores SNMP traps (SNMPV2c and SNMPV3) and information messages from devices. Routing Director listens for SNMP traps on the default UDP port 162. These traps are then forwarded to external managers (applications) that are configured in Routing Director. Routing Director, by default, stores SNMP Traps for one month.

Routing Director logs all the activities related to SNMP processing if the SNMP trap service (trapster) log level is configured as info. To view these logs, log into a Routing Director cluster node and execute the following command:

```
root@controller:~# kubectl logs -f -n trapster svc/trapster
```

To configure Routing Director to accept SNMP traps and export them to external SNMP managers:

1. Click **Settings Menu > System Settings** on the banner.

The Organization Settings page appears.

2. Configure SNMP parameters under **SNMP Forwarding Endpoints** and **SNMP Configuration**.

Refer to [Table 30 on page 78](#). Fields marked * are mandatory.

3. Click **Save** to save the SNMP configuration.

Table 30: SNMP Configuration Parameters

Field	Description
SNMP Forwarding Endpoints	
Name*	Enter a name for the external SNMP manager (application) to which SNMP traps received from Juniper devices must be sent. The SNMP manager name can include any type of character but must not exceed the 128-character limit.
Target Address*	Enter an IPv4 address, an IPv6 address, or the fully qualified domain name of the external SNMP manager. The Target Address can include any type of character but must not the exceed 255 characters.
Port	Enter the port number on the external SNMP manager to which Routing Director must send SNMP traps received from Juniper devices. Range: 0 through 65,535

Table 30: SNMP Configuration Parameters *(Continued)*

Field	Description
Description	Enter a description of the external SNMP manager.
SNMP Configuration V2	
Timeout (in seconds)*	Enter the maximum duration (in seconds) that Routing Director can take to forward SNMP traps received from Juniper devices to external SNMP managers. Range: 5 through 30 seconds
Retry*	Enter the number of times Routing Director can attempt forwarding an SNMP trap to external SNMP managers. Range: 0 through 5
Validate	Click to enable (default) or disable Routing Director from validating the configured community string with the community string in the SNMP trap received from Juniper devices, for authentication.
Community	Enter a string that Routing Director can use to verify the SNMP traps received from Juniper devices. If the community string you enter here matches with the community string in the received SNMP traps, Routing Director forwards the SNMP traps to the configured external SNMP managers. The community string must not exceed 32 characters.
SNMP Configuration V3	
Timeout (in seconds)*	Enter the maximum duration (in seconds) that Routing Director can take to forward SNMP traps received from Juniper devices, to external SNMP managers. Range: 5 through 30 seconds

Table 30: SNMP Configuration Parameters *(Continued)*

Field	Description
Retry*	Enter the number of times Routing Director can attempt forwarding an SNMP trap to external SNMP managers. Range: 0 through 5
Authentication Type	Select an algorithm (such as MD5 or SHA) that Routing Director can use to verify the identity of the SNMP user configured on Juniper devices sending the SNMP traps. The authentication type you select here must match the authentication type configured on the Juniper devices. If authentication fails, Routing Director will not forward the traps to external SNMP managers.
Username	Enter the SNMP username that Routing Director can use to verify the integrity of the SNMP traps received from Juniper devices. The username is used together with the authentication type and password to verify the SNMP user's identity configured on Juniper devices. The username you enter here must match the username configured on the Juniper devices. The username can contain any type of character but must not exceed 32 characters.
Authentication Password	Enter a secret key to be used together with the authentication type and username to verify the integrity of SNMP messages. The password is used together with the authentication type and username to verify the user's identity. The password you enter here must match the password configured on the devices. The password can contain any type of character but must not exceed the 32-character limit.
Encryption Type*	Select an algorithm, such as DES or AES, that Routing Director can use to decrypt SNMP traps to validate them and then encrypt the SNMP traps back before sending them to external SNMP managers.

Table 30: SNMP Configuration Parameters *(Continued)*

Field	Description
Privacy Password	Enter the password that Routing Director can use to decrypt and encrypt SNMP traps to validate them before forwarding to external SNMP managers. The privacy password can contain any type of character but must not exceed the 32-character limit.

RELATED DOCUMENTATION

[Manage Organization Settings](#) | 60

[Administration Workflow](#) | 55

Configure Webhooks to Receive Event Notifications

You use webhooks to automate sending event notifications from a source application to a destination application. You can configure webhooks to enable Routing Director to send notifications to third-party applications, such as Slack, when the events that you have subscribed to are triggered on the managed devices.



NOTE: Internet connectivity is required for Routing Director to connect to third-party applications.

The following section describes how to configure webhooks to receive event notifications in Slack.

To receive webhook notifications in a format that is compatible with Slack, you need to configure an intermediary that can interact with the sending and receiving applications, in this case, Routing Director and Slack. The recommended intermediary platform is Make. To process notifications, Make uses a workflow called Scenario, which converts the notifications to a format that Slack supports. Each event notification is sent to a URL that is generated for the Scenario in Make. The notification is then converted into a format that Slack supports and delivered to the configured Slack channel.

For information about Scenario in Make, see [Scenario](#).

To configure webhooks in Routing Director to send notifications to a Slack channel:

1. Log in to Make <https://www.make.com/en/login>. From the home page, navigate to Scenario on the left navigation menu.
2. Configure the scenario settings as described, see [Creating a Scenario](#).
Make generates a URL. Whenever an event is triggered, Routing Director sends webhook notifications to this URL.
3. In the Routing Director GUI, navigate to Organization Settings (**Settings Menu > System Settings**).
The Organization Settings page appears.
4. In the Webhooks tile, enable Webhooks.
5. Configure the webhooks settings. See [Table 31 on page 83](#) for webhooks field descriptions.

**NOTE:**

- In the URL field, enter the URL generated in step 2.
- Fields marked * are mandatory.

6. (Optional) Verify Webhook-Slack integration by logging in to the CLI of a device and generating an event.

For example, run the following commands in the device CLI to generate an alert.

```
user@host# set interfaces et-0/0/1 disable
user@host# commit
user@host# run show interfaces terse | grep et-0/0/1
et-0/0/1          down  down
user@host# delete interfaces et-0/0/1 disable
user@host# commit user@host# run show interfaces terse | grep et-0/0/1
et-0/0/1          up    down
```

7. (Optional) Verify that:

- The event you generated is listed on the Events page (**Observability > Events**).
- You received a notification for the event in the Slack channel.

**NOTE:**

- You must have access to the Slack channel to view event notifications in Slack.
- You must be an administrator with the Network Admin role to perform corrective action.

Table 31: Parameters to Configure Webhooks

Field	Description
Name*	Enter a name for the webhook. The name can contain alphanumeric and special characters.
URL*	Enter the URL generated in Make for the scenario.
Verify SSL Cert	Click to enable or disable (default) verifying SSL certificates from the device from which notifications are received.
Secret	Enter the secret to validate that the notifications received are from valid hosts. The secret can contain a string of alphanumeric and special characters.
Webhook Header	Webhook custom headers are key-value pairs that provide additional information about the notifications. You can add multiple custom headers to: • Provide additional information in plain text, along with the default headers, about the webhook notifications being sent to the configured endpoint. • Provide security, such as API keys, to verify end-to-end data integrity, for authorization, and so on. Click the Add icon (+) to add webhook headers. The Webhook Header page appears. • Header Key*—Enter a unique key. • Header Value*—Enter a unique value for the key. The value can contain alphanumeric characters. Click the Delete icon (trash can) to remove the webhook headers.

Table 31: Parameters to Configure Webhooks (*Continued*)

Field	Description
Streaming APIs	Enable the events for which you want to receive notifications. You can subscribe to events such as, alerts, audits, device status, and device alarms to get real-time notifications when the event occurs. • Alerts—Click to enable or disable receiving notifications when subscribed alerts are generated on the managed devices. Alerts notification is disabled by default. • Audits—Click to enable or disable receiving notifications when a user accesses an organization or modifies organization settings. Audits notification is disabled by default. • Device Status—Click to enable or disable receiving notifications when the device status changes due to events such as a link going up or down, or the device getting disconnected from Routing Director, and so on. The Device Status notification is disabled by default. • Device Alarms—Click to enable or disable receiving notifications when subscribed alarms are generated on the managed devices. Device Alarm notification is disabled by default.

RELATED DOCUMENTATION

[Manage Organization Settings](#) | 60

[Administration Workflow](#) | 55

Manage RADIUS Server Configurations

SUMMARY

Use this topic to add, edit, or delete RADIUS Server configurations.

IN THIS SECTION

- [Add a RADIUS Server](#) | 85

Add a RADIUS Server

IN THIS SECTION

- [Edit a RADIUS Server](#) | 86

To configure a RADIUS Server in Routing Director:

1. Click **Settings menu > System Settings** on the banner.
The Organization Settings page appears.
2. Under Device Authentication, click **Radius**.
The Radius Configuration table appears.
3. Click the **+** (Add) icon to configure a RADIUS Server.
The Create Radius Configuration page appears.
4. Enter details by referring to [Table 32 on page 85](#).
Fields marked * are mandatory.
5. Click **Save**.
The RADIUS Server is configured and listed in the RADIUS Configuration table.

Table 32: Parameters to Configure a RADIUS Server

Field	Description
Server Address*	Enter the IP address (IPv4 or IPv6) of the server.
Secret*	Enter the password that the network device would use to authenticate with the RADIUS Server.
Port	Enter the port on which to contact the RADIUS Server. Default: 1812 Range: 0 through 65,535

Table 32: Parameters to Configure a RADIUS Server *(Continued)*

Field	Description
Accounting Port	Enter the port to be used for accounting purposes. Default: 1813 Range: 0 through 65,535
Retry	Enter the number of times that the device attempts to contact the RADIUS Server. If you enter a value outside the range, you are prompted to re-enter the value. Range: 1 through 100
Timeout	Enter the time that the device can wait (in seconds) to receive a response from the RADIUS Server. If you enter a value outside the range, you are prompted to re-enter the value again. Range: 1 through 100
Username*	Enter the username that would be used to onboard the device. There are no restrictions on the number or type of characters that you can use.
Password*	Enter the password for the username. There are no restrictions on the number or type of characters that you can use.

Edit a RADIUS Server

IN THIS SECTION

- [Delete a RADIUS Server | 87](#)

To edit a RADIUS Server:



NOTE: The changes to the RADIUS Server configuration are automatically saved to the database. However, you must click **Save** to save the change in device authentication type.

1. Click **Settings menu > System Settings** on the banner.

The Organization Settings page appears.

2. Click the RADIUS Server you want to edit.

The Edit RADIUS Configuration page appears.

3. Edit the RADIUS Server configuration by referring to [Table 32 on page 85](#).

4. Click **Save**.

You are returned to the Organization Settings page, where you can view the changes in the RADIUS Configuration table.

Delete a RADIUS Server

To delete a RADIUS Server:

1. Click **Settings menu > System Settings** on the banner.

The Organization Settings page appears.

2. Click the RADIUS Server you want to edit.

The Edit RADIUS Configuration page appears.

3. Click **Delete**.

You are returned to the Organization Settings page, where you can view that the RADIUS Server configuration is removed from the RADIUS Configuration table.

SEE ALSO

[Manage Organization Settings | 60](#)

[Administration Workflow | 55](#)

CHAPTER 6

Site Management

IN THIS CHAPTER

- [About the Sites Page | 88](#)
- [Add Sites | 90](#)
- [Edit and Delete Sites | 91](#)

About the Sites Page

IN THIS SECTION

- [Tasks You Can Perform | 88](#)
- [Field Description | 89](#)

Sites are the physical locations that host devices, such as routers, switches, and firewalls within an organization's network. The superuser can create sites and add devices to those sites. Sites are used to identify the location of the devices in the organization. Multiple sites can be grouped into site groups for easy management. For more information on organizations and sites, see ["Organization and Sites Overview" on page 57](#).

To access the Sites page, click **Inventory > Common Resources > Sites**.

Tasks You Can Perform

You can perform the following tasks from this page:

- View details about the sites in an organization—You can view the site name, country, time zone, address, the site group to which the site belongs, and notes about the site.
- Add sites; see ["Add Sites" on page 90](#).

- Modify and delete sites; see ["Edit and Delete Sites" on page 91](#).
- Filter the data displayed in the table—Click the filter icon (funnel) and select whether you want to show or hide advanced filters. You can then add or remove filter criteria, save criteria as a filter, apply or clear filters, and so on. The filtered results are displayed on the same page.
- Search by using keywords—Click the search icon (magnifying glass), enter the search term in the text box, and press Enter. The search results are displayed on the same page.
- Show or hide columns in the table or reset page preferences, using the vertical ellipsis menu.
- Sort, resize, or re-arrange columns in a table (grid).
- Set the number of rows to be displayed in the table—Set the number of rows (15, 30, 50, or 100) to be displayed in the Network Implementation Plan table in the **Display** drop-down list provided at the bottom right-corner of the table.

If there are multiple pages of rows, use the arrow heads to traverse through the pages or enter a page number to view the entries listed on that page.

- Refresh the Sites table—Click the **Refresh** icon present at the bottom of the table to view the most recent updates made to sites.

Field Description

[Table 33 on page 89](#) describes the fields displayed on the Sites page.

Table 33: Fields on the Sites Page

Fields	Description
ID	ID of the site.
Name	Displays the name of the site.
Country	Displays the country where the site is located.
Timezone	Displays the time zone of the site.
Address	Displays the address of the site.
Site Groups	Displays the site groups to which the site belongs, if any.
Notes	Displays additional information about the site.

RELATED DOCUMENTATION

About the Inventory Page 128
Organization and Sites Overview 57

Add Sites

A site identifies the location of the devices in an organization. A superuser can add sites in an organization.

To add a site:

1. Click **Inventory > Common Resources > Sites** in the left navigation menu.
The Sites page appears.
2. Click **Create Site (+)** icon.
The Create Site page appears.
3. Enter the site parameters, select a valid location, and site groups according to the guidelines provided in [Table 34 on page 90](#).
Fields marked * are mandatory.
4. Click **Save**.
A confirmation message indicating that the site is created is displayed, and the site is listed on the Sites page.

Table 34: Fields on the Create Site Page

Fields	Description
Name*	Enter a unique name for the site. The site name can contain up to 64 characters.
Country*	Select the country where the site is located. If you enter the coordinates or enter a location in the Location Search field, Country is updated with the respective country. However, if you select a country from the drop-down list, the country is not reflected on the map.
Timezone*	Select the timezone of the site. If you enter coordinates or enter a location in the Location Search field, Timezone is updated with the respective timezone. However, if you select a country from the drop-down list, the timezone is not reflected on the map.

Table 34: Fields on the Create Site Page *(Continued)*

Fields	Description
Location*	Specify the location of the site. The Country and the Timezone fields for the site are automatically updated when you perform any of the following: • Click the location of the site on the map. • Enter the coordinates in the Search field. • Enter the location in the Search field.
Site Groups	Select the site groups, if any, to which the site should belong. If a site group is not available, you can type a name for the site group and press Enter to create the site group.
Notes	Enter additional information about the site. Notes can contain up to 1000 characters.

RELATED DOCUMENTATION

[Organization and Sites Overview](#) | 57

[Assign a Device to a Site](#) | 133

Edit and Delete Sites

IN THIS SECTION

● [Edit Site Information](#) | 92

● [Delete a Site](#) | 92

A superuser can modify site information and delete sites in an organization.

Edit Site Information

To modify site details:

1. Click **Inventory > Common Resources > Sites** in the left navigation menu.

The Sites page appears.

2. Select the site whose details you want to edit and click **Edit Site** (pencil) icon.

The Edit Site: *Site Name* page appears.

3. Modify the site details as needed.

See [Table 34 on page 90](#) for field descriptions.

4. Click **Save**.

A confirmation message indicating that the site details are updated is displayed and you are returned to the Sites page, where you can verify the changes you made.

Delete a Site

You can delete sites from the organization.



NOTE: When you delete a site, the site is removed permanently from the organization.

To delete a site:

1. Click **Inventory > Common Resources > Sites** in the left navigation menu.

The Sites page appears.

2. Select the site you want to delete and click **Delete Site** (trash) icon.

The Delete Site confirmation page appears.

3. Click **OK**.

A confirmation message indicating that the site is deleted is displayed and you are returned to the Sites page.

SEE ALSO

[Organization and Sites Overview | 57](#)

[About the Sites Page | 88](#)

CHAPTER 7

User Management

IN THIS CHAPTER

- [About the Users Page | 93](#)
- [Predefined User Roles Overview | 97](#)
- [Configure SMTP Settings in Paragon Shell | 101](#)
- [Add Users to an Organization Overview | 107](#)
- [Invite Users | 109](#)
- [Manage Users and Invites | 112](#)
- [Manage Paragon Shell Users | 117](#)
- [Manage Your Routing Director Account | 125](#)

About the Users Page

IN THIS SECTION

- [Tasks You Can Perform | 94](#)
- [Field Descriptions | 95](#)

To access the Users page, click **Settings Menu > Users** on the banner.

You must be an administrator with the Super User role to view, add, and manage users. Depending on whether SMTP is configured or not, different workflows exist to add and manage users in Routing Director. See "[Configure SMTP Settings in Paragon Shell](#)" on page 101.

If SMTP is configured, Routing Director sends an invitation e-mail when:

- A first-time user creates an account in Routing Director

- A superuser adds a user to an organization.

If SMTP is not configured in Routing Director, an administrator with the Super User role creates a user account in the organization and shares the Routing Director URL, e-mail ID, and temporary password with the user. The user must log in with the credentials provided by the superuser to activate the account. Once logged into the organization, the user is prompted to create a new password. For more information on user login, see ["User Activation and Login" on page 39](#).

Tasks You Can Perform

An administrator with the Super User role can perform the following tasks from this page:

- View details of the existing users and the users who are invited to access the organization—The basic information about the users, such as first name, last name, e-mail ID, invite status of the user, and role assigned is displayed. See [Table 35 on page 95](#) for field descriptions.
- Invite users; see ["Invite Users" on page 109](#).
- Manage user invitations; see ["Manage Users and Invites" on page 112](#).
- Assign tags to users—A tag is a label in the key:value format that you add to a user so that you can logically group and categorize the users.

To add one or more tags to a user, on the Users page:

1. Select the user and click the **Edit** icon.

The Edit Users page appears.

2. In the **Tags** field, select a tag to be applied from the drop-down list that appears.

Alternatively, you can also create a new tag by clicking the **Create New Tag** link at the bottom of the page and assign it to the user. If you are creating a new tag, ensure that:

- The key starts with a letter and is followed by letters, numbers, periods (.), hyphens (-), or underscores (_), and is between 1 and 200 characters long.
- The value includes alphanumeric characters, periods (.), hyphens (-), and underscores (_), and is between 1 and 200 characters long.

3. Click **Save** to save the changes.

- Resize or re-arrange columns in a table (grid).
- Show or hide columns in the table or reset page preferences, using the vertical ellipsis menu.
- Search by using keywords—Click the search icon (magnifying glass), enter the search term in the text box, and press Enter. The search results are displayed on the same page.

- Filter the data displayed in the table—Click the filter icon (funnel) and select whether you want to show or hide advanced filters. You can then add or remove filter criteria, save criteria as a filter, apply or clear filters, and so on. The filtered results are displayed on the same page.
- Set the number of rows to be displayed in the table—Set the number of rows (15, 30, 50, or 100) to be displayed in the Network Implementation Plan table in the **Display** drop-down list provided at the bottom right-corner of the table.

If there are multiple pages of rows, use the arrow heads to traverse through the pages or enter a page number to view the entries listed on that page.

For more information, see ["GUI Overview" on page 6](#).

Field Descriptions

[Table 35 on page 95](#) describes the fields on the Users page.

Table 35: Fields on the Users Page

Fields	Description
First Name	The first name of the user.
Last Name	The last name of the user.
Email	The e-mail ID the user would use to access Routing Director.

Table 35: Fields on the Users Page *(Continued)*

Fields	Description
Status	Indicates a user's account status: User account status when SMTP is configured: • Active: The user's account is active and the user can access the organization. • Invite Pending: The user is yet to accept the e-mail invitation received and doesn't have access to the organization. • Invite Expired: The e-mail invitation sent to the user has expired. An invitation expires seven days after the invite is sent to the user. User account status when SMTP is not configured: • Created: The superuser has created a user account in Routing Director for the user. The user must login using the e-mail ID and temporary password that the superuser shares with them to activate their account. • Active: The user's account is active and the user can access the organization.
Role	The role assigned to a user. See "Predefined User Roles Overview" on page 97 for details about the user roles.
Tags	Tags assigned to the user. If there is more than one tag assigned to the user, an integer is present next to the displayed tag. For example, +2 in the Tags column indicates that there are two more tags in addition to what is currently listed. You can hover over the integer to view the additional tags assigned to the user. To assign a tag, edit the user and click Edit. For information, see "Manage Users and Invites" on page 112

RELATED DOCUMENTATION

| [Organization and Sites Overview](#) | 57

Predefined User Roles Overview

Routing Director provides four predefined roles to manage access privileges of users, based on the tasks users need to perform.

A superuser creates an organization, adds users to predefined roles depending on the requirements of the organization. For example, an organization with a large number of networking devices would require multiple users performing different roles to efficiently manage the organization, whereas, in a small organization, a single user can perform all the four roles. Different types of users in an organization, such as a network architect, network planner, NOC engineer, and field technician, all derive their access privileges from the predefined roles assigned to them.

User Roles and their Responsibilities

The four predefined roles are:

- Super User
 - Is the administrator of the organization.
 - Creates organization, invites users, assigns user roles, creates sites, adopts devices, and so on.
 - A superuser doesn't need to be a person with a high-level of networking domain expertise.
- Network Admin
 - Is a networking expert who monitors, verifies, and troubleshoots an organization's network.
- Observer
 - Monitors events in the organization's network.
 - An Observer cannot take corrective action. The observer brings issues to the notice of the network administrator for resolution.
- Installer
 - Onboards devices and monitors device status during onboarding.
 - An Installer can access only the Onboard a Device and Device List pages.

[Table 36 on page 98](#) displays the access privileges of the four user roles to the menu items.

Table 36: User roles and their access privileges

Menu	Super User	Network Admin	Observer	Installer
Observability				
Observability > Health				
Troubleshoot Devices	✓	✓	✓	✗
Events	✓	✓	✓	✗
Observability > Active Assurance				
Monitors	✓	✓	✓	✗
Tests	✓	✓	✓	✗
Measurement Designer	✓	✓	✓	✗
Observability > Network				
Topology	✓	✓	✓	✗
Trust and Compliance				
Trust > General				
Trust Plans	✓	✓	✓	✗
Trust > Compliance				
Compliance Scan	✓	✓	✓	✗
Checklist	✓	✓	✓	✗

Table 36: User roles and their access privileges *(Continued)*

Menu	Super User	Network Admin	Observer	Installer
Benchmarks	✓	✓	✓	✗
Tailorings	✓	✓	✓	✗
Trust > Integrity				
Hardware EOL	✓	✓	✓	✗
Software EOL	✓	✓	✓	✗
Trust > Vulnerabilities				
Advisories	✓	✓	✓	✗
PBNs	✓	✓	✓	✗
Orchestration				
Orchestration > Service				
Instances	✓	✓	✓	✗
Service Catalog	✓	✓	✓	✗
Customers	✓	✓	✓	✗
Orchestration > Monitoring				
Workflows	✓	✓	✓	✗
Inventory				

Table 36: User roles and their access privileges *(Continued)*

Menu	Super User	Network Admin	Observer	Installer
Inventory > Devices				
Network Inventory	✓	✓	✓	✗
Device and Interface Profiles	✓	✓	✓	✗
Configuration Templates	✓	✓	✓	✗
Configuration Backups	✓	✓	✓	✗
Software Images	✓	✓	✓	✗
Inventory > Device Onboarding				
Onboarding Dashboard	✓	✓	✓	✗
Network Implementation Plan	✓	✓	✓	✗
Inventory > Active Assurance				
Test Agents	✓	✓	✓	✗
Inventory > Common Resources				
Sites	✓	✗	✗	✗
Tags	✓	✗	✗	✗
Settings Menu				

Table 36: User roles and their access privileges *(Continued)*

Menu	Super User	Network Admin	Observer	Installer
System Settings	✓	×	×	×
Users	✓	×	×	×
Audit Logs	✓	✓	×	×
Field Technician UI				
Onboard a Device	×	×	×	✓
Device List	×	×	×	✓

RELATED DOCUMENTATION
[Add Access Control Profiles | 154](#)
[Manage Identity Providers | 68](#)
Configure SMTP Settings in Paragon Shell**IN THIS SECTION**

- [Enable SMTP Based User Management | 102](#)
- [Disable SMTP Based User Management | 105](#)

SMTP must be configured in Routing Director for Routing Director to send e-mail notifications to users when their account is created or modified and when subscribed events occur. By default, SMTP is not configured in Routing Director and notifications are not sent.

Table 37 on page 102 lists the behavior for different scenarios when SMTP and local authentication are configured.

Table 37: Behavior for Different Scenarios when SMTP and Local Authentication are configured

Local Authentication	SMTP	Behavior
Enabled	Disabled	User is authenticated locally based on user accounts created within Routing Director. E-mail notifications are not sent to users.
Enabled	Enabled	User is authenticated locally based on user accounts created within Routing Director. E-mail notifications are sent when subscribed alerts or events occur.
Disabled	Enabled	Users are invited to access an organization over an e-mail. Also, users receive notifications for subscribed alerts or events over an e-mail.
Disabled	Disabled	This is not a valid combination. However, a user is locally authenticated based on user account created in Routing Director. E-mail notifications are not sent for subscribed alerts or events.

A system administrator can configure SMTP in Paragon Shell so that users can be notified when their account is created, activated, or modified.

Enable SMTP Based User Management

To enable SMTP settings in Paragon Shell:

1. SSH to the node on which you deployed the Routing Director cluster.
2. Log in to Paragon Shell as the root user.
You are placed in the operational mode.
3. Enter the configuration mode.

```
root@eop-primary> configure
Entering configuration mode
```

```
[edit]
```

4. Configure the following SMTP parameters.

```
root@eop-primary# set paragon cluster mail-server smtp-allowed-sender-domains sender-domains
```

```
[edit]
```

```
root@eop-primary# set paragon cluster mail-server smtp-relayhost relayhost-hostname
```

```
[edit]
```

```
root@eop-primary# set paragon cluster mail-server smtp-relayhost-username relayhost-username
```

```
[edit]
```

```
root@eop-primary# set paragon cluster mail-server smtp-relayhost-password relayhost-password
```

```
[edit]
```

```
root@eop-primary# set paragon cluster mail-server smtp-sender-email sender-e-mail-address
```

```
[edit]
```

```
root@eop-primary# set paragon cluster mail-server smtp-sender-name sender-name
```

```
[edit]
```

```
root@eop-primary# set paragon cluster mail-server smtp-enabled true
```

Where:

smtp-allowed-sender-domains are the e-mail domains from which Routing Director sends e-mails to users.

relayhost-hostname is the name of the SMTP server that relays messages.

relayhost-username (optional) is the username to access the SMTP (relay) server.

relayhost-password (optional) is the password for the SMTP (relay) server.

sender-e-mail-address is the e-mail address that appears as the sender's e-mail address to the e-mail recipient.

sender-name is the name that appears as the sender's name in the e-mails sent to users from Routing Director.

5. Disable local user management to perform SMTP-based user management configurations.

```
root@eop-primary# set paragon cluster papi papi-local-user-management false
```

6. Commit the changes and exit the configuration mode.

```
root@eop-primary# commit and-quit
commit complete
Exiting configuration mode
```

7. Add the SMTP configuration parameters to the configuration files.



NOTE: Ensure that the following parameters are configured in the **config.yml** configuration file before adding the SMTP parameters to the file.

- ntp_servers
- ingress_vip
- test_agent_gateway_vip
- web_admin_user
- web_admin_password

```
root@eop-primary> request paragon config
Paragon inventory file saved at /epic/config/inventory
Paragon config file saved at /epic/config/config.yml
```

8. Deploy the SMTP configuration on the node.

```
root@eop-primary> request paragon deploy cluster input "-t addon-apps,papi -e
target_components=papi-kafka-emitter,mailservice"
Process running with PID: 2xxxxx2
```

The deployment takes about two to three minutes to complete.

9. (Optional) Monitor the progress of the deployment.

```
root@eop-primary> monitor start /epic/config/log
```

After SMTP is enabled, when a superuser invites users to an organization, Routing Director sends e-mail invitations to the users. Routing Director also validates user accounts and notifies users of any modifications to their accounts through e-mails.

Disable SMTP Based User Management

To disable SMTP settings in Paragon Shell:

1. SSH to the node on which you deployed the Routing Director cluster.
2. Log in to Paragon Shell as the root user.
You are placed in the operational mode.
3. Enter the configuration mode.

```
root@eop-primary> configure
Entering configuration mode

[edit]
```

4. Disable the SMTP-based user management configurations.

```
root@eop-primary# set paragon cluster mail-server smtp-enabled false
```

5. Enable local user management.

```
root@eop-primary# set paragon cluster papi papi-local-user-management true
```

6. Commit the changes and exit the configuration mode.

```
root@eop-primary# commit and-quit
commit complete
Exiting configuration mode
```

7. Add the updated configuration parameters to the configuration files.



NOTE: Ensure that the following parameters are configured in the **config.yml** configuration file before adding the SMTP parameters to the file.

- ntp_servers
- ingress_vip
- test_agent_gateway_vip
- web_admin_user
- web_admin_password

```
root@eop-primary> request paragon config
Paragon inventory file saved at /epic/config/inventory
Paragon config file saved at /epic/config/config.yml
```

8. Deploy the SMTP configuration on the node.

```
root@eop-primary> request paragon deploy cluster input "-t addon-apps,papi -e
target_components=papi-kafka-emitter,mailservice"
Process running with PID: 2xxxxx2
```

The deployment takes about two to three minutes to complete.

SMTP is disabled in Routing Director.

Users will not be notified when their account is created, activated, or modified. After adding users to an organization, superusers must manually share the Web URL of Routing Director and credentials (e-mail ID and temporary password) to log in, with the users.

RELATED DOCUMENTATION

[Manage Organization Settings](#) | 60

Add Users to an Organization Overview

IN THIS SECTION

- [Add Users to an Organization When SMTP is Configured | 107](#)
- [Add Users to an Organization When SMTP is Not Configured | 108](#)

An administrator with the Super User role can add or invite users to an organization from the Routing Director GUI. Depending on whether SMTP is configured or not, two different workflows exist for adding users to an organization in Routing Director.

The workflows are as follows:

- ["Add Users to an Organization When SMTP is Configured" on page 107](#)
- ["Add Users to an Organization When SMTP is Not Configured" on page 108](#)


Existing users can access their organization by using their Routing Director account.

Add Users to an Organization When SMTP is Configured

When SMTP is configured in Routing Director, an administrator with the Super User role can invite users to an organization and provide role-based access by sending an invitation to the user's e-mail ID. The user needs to accept the invitation to be a member of the organization.

[Figure 10 on page 108](#) illustrates the workflow for inviting a new user to an organization when SMTP is configured.

Figure 10: Add users to an organization when SMTP is configured



The status of the invitation is shown as Invite Pending until the user:

- Accepts the invitation.
- Rejects the invitation.

If the user doesn't accept or reject the invitation within seven days. The status of the invitations changes to Invite Expired.

If the user accepts the invitation and has role-based access to the organization, but you want to take away the user's access, you can revoke the invitation. See ["Revoke a User" on page 115](#).

If the user invitation expires, the superuser can re-invite the user or cancel the invitation.

Add Users to an Organization When SMTP is Not Configured

When SMTP is not configured in Routing Director, the superuser adds users to an organization and manually shares the Routing Director Web URL and credentials (e-mail address and temporary password) to log in, with the users. See ["Invite Users When SMTP is Disabled and Local Authentication is Enabled" on page 110](#).

Users can access the Web URL and enter the e-mail address and temporary password to log in to Routing Director. On logging in, the user is prompted to set a new password. See ["Log in as a New User Without an Invite \(when SMTP is not Configured\)" on page 40](#).

The status of the user account is shown as Created until the user logs in to Routing Director and accesses the organization to which they were added. On successful login, the status of the account becomes Active.

RELATED DOCUMENTATION

[Authentication Methods Overview | 65](#)

[Organization and Sites Overview | 57](#)

Invite Users

IN THIS SECTION

- [Invite Users when SMTP is Enabled | 110](#)
- [Invite Users When SMTP is Disabled and Local Authentication is Enabled | 110](#)

An administrator with the Super User role can add users to an organization from the Routing Director GUI. Depending on whether SMTP is configured or not, two different workflows exist for adding users to an organization in Routing Director. See ["Configure SMTP Settings in Paragon Shell" on page 101](#).

To invite users to an organization in Routing Director when SMTP is configured, see ["Invite Users when SMTP is Enabled" on page 110](#). The user must accept the invitation within seven days, after which the invitation expires.



NOTE: A new user without an invite can create an account from the login page of the Routing Director Web GUI, when SMTP is configured. See, ["Log in as a New User for the First Time Without an Invite \(when SMTP is Configured\)" on page 41](#).

To add users to an organization in Routing Director when SMTP is not configured, see ["Invite Users When SMTP is Disabled and Local Authentication is Enabled" on page 110](#).

Invite Users when SMTP is Enabled

To invite a user when SMTP is configured:

1. Click **Settings Menu > Users** on the banner.
The Users page appears.
2. Click the **Invite User (+)** icon.
The Users: New Invite page appears.
3. Enter user details and assign a role according to the guidelines provided in [Table 38 on page 111](#).
Fields marked * are mandatory.
4. Click **Invite**.
A confirmation message indicating that the user is invited is displayed, and the user details are listed on the Users page.
5. Check the status of the user. If the status changes to Invite Expired, you can delete the user, reinvite the user, or cancel the invitation. For more information, see ["Manage Users and Invites" on page 112](#).

Invite Users When SMTP is Disabled and Local Authentication is Enabled

To invite a user when SMTP is disabled and local authentication is enabled:

1. Click **Settings Menu > Users** on the banner.
The Users page appears.
2. Click the **Invite User (+)** icon.
The New User page appears.
3. Enter user details and assign a role according to the guidelines provided in [Table 38 on page 111](#).



NOTE:

- Each user must have a unique and valid e-mail address as their username.
- Fields marked * are mandatory.

4. Click **Save**.
The New User Creation window appears. You are presented with a masked temporary password.

You must manually share this password with the user for them to access the organization. For more information on user login, see ["User Activation and Login" on page 39](#).
5. Click **Copy to Clipboard** to copy the masked password.
To view the masked password, click **Show Password**.
6. Click **OK**.

The user is listed on the Users page with the status as Created. The status changes to Active when the user successfully logs in to the application.

7. In case, you choose to not share the password at the time of adding the user, you can later select the user from the users table, click **Edit > Reset Password** to regenerate a new temporary password and share it with the user.

Table 38: Fields on the Invite User Page

Field	Description
First Name*	Enter the first name of the user. First name can contain up to 64 characters.
Surname*	Enter the last name of the user. Last name can contain up to 64 characters.
Email*	Enter the e-mail ID that a user would use to access Routing Director.
Role*	Assign a role to the user. You can assign only one role to a user in an organization. You can assign: • Super User • Network Admin You must assign at least one access control profile to a Network Admin to grant access to resources. See "Access Control Profile Overview" on page 145. • Observer • Installer See "Predefined User Roles Overview" on page 97 for information about privileges associated with each role.
Access Control Profile*	Assign one or more access control profiles to grant access to resources (devices, service instances, and perform tagging operations). If an access profile does not exist, click the +Add Access Control Profile link to create a new profile and assign it to the Network Admin user. See "Add Access Control Profiles" on page 154.

SEE ALSO

| [Reset Your Password](#) | 42

Manage Users and Invites

IN THIS SECTION

- [Edit User Information when SMTP is Configured](#) | 113
- [Reinvite a User](#) | 113
- [Cancel an Invitation](#) | 114
- [Revoke a User](#) | 115
- [Edit User Information when SMTP is not Configured](#) | 115
- [Reset Your Password](#) | 116
- [Delete a User](#) | 116

You must be a superuser to manage users. Depending on whether SMTP is configured or not, different options exist for managing users in Routing Director.

If you configure SMTP in Routing Director, the user is notified when they are invited to access an organization and when the user invite is modified. See "[Configure SMTP Settings in Paragon Shell](#)" on [page 101](#).

As a superuser, you can:

- Edit user information . See "[Edit User Information when SMTP is Configured](#)" on [page 113](#)
- Reinvite a user. See "[Reinvite a User](#)" on [page 113](#)
- Cancel invitations. See "[Cancel an Invitation](#)" on [page 114](#)
- Revoke a user. See "[Revoke a User](#)" on [page 115](#)
- Delete a user. See "[Delete a User](#)" on [page 116](#)

If SMTP is not configured in Routing Director, the superuser creates a user account in an organization and shares the Routing Director URL, and credentials (e-mail ID, and temporary password) with the user. The user must log in to Routing Director with the credentials provided by the superuser to activate the

account. Once logged into the organization, the user is prompted to create a new password. For more information, see ["User Activation and Login" on page 39](#).

As a superuser, you can:

- Edit user information. See ["Edit User Information when SMTP is not Configured" on page 115](#)
- Reset the user's password. See ["Reset Your Password" on page 116](#)
- Delete user. See ["Delete a User" on page 116](#)

Edit User Information when SMTP is Configured

On the User: *User-Name* page, you can assign a pre-defined tag or edit the role and tags assigned to a user. If a tag does not exist, you can create a tag by clicking the **Create New Tag** link and then assign the tag to the user

You cannot modify the first name, last name, and e-mail ID of a user.

To edit a user's role:

1. Click **Settings Menu > Users** on the banner.

The Users page appears.

2. Select the user whose role you want to edit and click **Edit User** (pencil) icon.

The User: *User-Name* page appears.

3. Assign a pre-defined tag or modify the role or assigned tags as needed. See [Table 35 on page 95](#) for field descriptions.



NOTE:

- If you modify the role or the tags assigned to a user whose invitation status is Active, the user is not notified about the modification in the role.
- If you modify the role or the tags assigned to a user whose invitation status is Invite Pending or Invite Expired, a new invitation e-mail is sent to the user to access the organization with the new role-based access privileges.

4. Click **Save**.

A confirmation message indicating that the user invitation is updated is displayed and you are returned to the Users page, where you can view the changes you made.

Reinvite a User

When SMTP is configured, you can reinvite a user if:

- The user invitation has expired. An invitation is valid only for seven days.

- The user invitation is pending; for example, when the user has not accepted the invitation.
- The user role needs to be modified for users with Invite Pending or Invite Expired invitation status.

To reinvite a user to the organization:

1. Click **Settings Menu > Users on the banner.**

The Users page appears.

2. Select the user you want to reinvite and do one of the following:

- Click **Edit User** (pencil) icon and then click **Re-invite** on the User: *User-Name* page.

You can modify the role of a user only when you reinvite the user by using this method.

- Click **More > Re-invite User**.
- Right-click the user in the table and click **Re-invite User**.

If you choose to reinvite a user by right-clicking or clicking the More option, you will only be able to reinvite the user but won't be able to modify their role.

The Re-invite User confirmation page appears.

You can reinvite a user whose status is Invite Expired or Invite Pending. For users whose access is revoked or deleted, you must reinvite the user by clicking **Invite User (+)** icon; see "[Invite Users](#)" on [page 109](#).

3. Click **Save.**

An invitation e-mail is sent to the user and the user account is listed on the Users page with status Invite Pending.

If the user doesn't accept the invitation within seven days, the invitation expires.

Cancel an Invitation

When SMTP is configured, you can invalidate an invitation by canceling the invitation. You can uninvite a user if the invitation status is Invite Pending or Invite Expired on the Users page.



NOTE: An invite expires after seven days from the day the invitation was sent.

To uninvite a user:

1. Click **Settings Menu > Users on the banner.**

The Users page appears.

2. Select the user you want to uninvite and do one of the following:

- Click **Edit User** (pencil) icon and then click **Uninvite** on the User: *User-Name* page.

- Click **More > Uninvite**.
- Right-click the user in the table and click **Uninvite**.

The Delete Invitation confirmation page appears.

3. Click **OK** to uninvite the user.

A confirmation message indicating that the invite is canceled appears and you are returned to the Users page. The details about the user invitation is no longer listed in the Users table.

Revoke a User

If the user accepts the invitation and has role-based access to the organization, but you want to take away the user's access, you can revoke the invitation. Revoking a user's access deletes the user from the organization. The user can still log in to Routing Director, but the user cannot access the organization to which the user's access got revoked. You can revoke access only for active accounts when SMTP is configured in Routing Director.

To revoke a user's access to an organization:

1. Click **Settings Menu > Users** on the banner.

The Users page appears.

2. Select the user whose access needs to be revoked and do one of the following:

- Click **Edit User** (pencil) icon and then click **Revoke** on the User: *User-Name* page.
- Click **More > Revoke User**.
- Right-click the user in the table and click **Revoke User**.

The Delete User confirmation page appears.

3. Click **OK**.

The user is deleted from the organization and cannot access the organization.



NOTE: Routing Director maintains a log of the user's activities in the organization even after the user's access gets revoked. For example, the user's activities recorded in the audit logs will remain even if they no longer have access to the organization.

Edit User Information when SMTP is not Configured

On the User: *Name* page, you can edit the role and the tags assigned to a user. A superuser cannot modify the first name, last name, and e-mail ID of a user.

To edit a user's role:

1. Click **Settings Menu > Users** on the banner.

The Users page appears.

2. Select the user you want to edit and click **Edit User** (pencil) icon.

The User: *User-Name* page appears.

3. Add a pre-defined tag and modify the role or assigned tags as needed. If a tag does not exist, you can create a tag by clicking the **Create New Tag** link. See [Table 35 on page 95](#) for field descriptions.
4. Click **Save**.

The updated details of the user are listed on the Users page.

Reset Your Password

To reset the password for a user when SMTP is not configured:

1. Click **Settings Menu > Users** on the banner.

The Users page appears.

2. Select the user whose password you want to reset and do one of the following:



NOTE: You can reset the password of users whose status is Created or Active.

- Click **Edit User** (pencil) icon and then click **Reset Password** on the User: **User-Name** page.
- Click **More > Reset Password**.
- Right-click the user and click **Reset Password**.

The Reset Password page appears. You are presented with a randomly generated temporary password.

You must share this password with the user for the user to access the organization.

3. Click **Copy to Clipboard** to copy the masked password.

To view the masked password, click **Show Password**.

You must manually share this password with the user for them to access the organization.

Delete a User

To delete a user from an organization:

1. Click **Settings Menu > Users** on the banner.

The Users page appears.

2. Select the user you want to deny access to the organization and do one of the following:
 - Click **Edit User** (pencil) icon and then click **Delete** on the User: *User-Name* page.
 - Click **More > Delete User**.

- Right-click the user and click **Delete User**.

The Delete User confirmation page appears.

3. Click **OK**.

The user is deleted from the organization and cannot access the organization.

SEE ALSO

| [Access the Routing Director GUI | 38](#)

Manage Paragon Shell Users

IN THIS SECTION

- [Create a User | 117](#)
- [Modify User Information | 120](#)
- [Delete a User | 122](#)
- [Logging in to Paragon Shell as a New User | 123](#)
- [Retrieve User Information from a Recovered Node in the Routing Director Cluster | 124](#)

System administrators (root user) use Paragon Shell to add users so that these users can access and manage Routing Director cluster configurations. The users created in Paragon Shell and Routing Director Web GUI are not shared across the two user interfaces (UIs).

The following sections describe the user management tasks you can perform by using Paragon Shell.



NOTE: User information configured on one node is deployed across all the nodes in the Routing Director cluster.

Create a User

System administrators and superusers can use Paragon Shell to create users who can access and manage the Routing Director cluster based on the privileges defined by the role assigned to them.

To create a user:

1. SSH to a node in the Routing Director cluster.

2. Log in to Paragon Shell.

If you are a system administrator, log in as the root user. If you are a superuser, log in with your login credentials shared by the system administrator.

You are placed in the operational mode of Paragon Shell.

3. Enter the configuration mode.

```
root@eop-primary> configure
Entering configuration mode

[edit]
```

4. Configure the username, access privileges, and authentication method to authenticate the user.

```
root@eop-primary# set system login user username class class authentication authentication-
method
```

Where,

- *username* is the unique name that identifies the user.
- *class* is the access privilege assigned to the user. The options are:
 - *read-only*—User can access Paragon Shell and view details about the Routing Director cluster but cannot enter configuration mode to modify cluster configuration.
 - *super-user*—User can access Paragon Shell and enter configuration mode to modify cluster configuration.
- *authentication* is the possible authentication methods that can be used to authenticate users. The options are:
 - *plain-text-password*—Enter the password in plain text. The password can contain alphanumeric and special characters and must have a minimum of six characters. The password is stored in an encrypted format in the configuration file.
 - *encrypted-password*—Enter the encrypted password and enclose it in quotation marks. Currently, the supported encryption algorithm is SHA-512.
 - *ssh-algorithm*—Generate an SSH keypair and use the key to authenticate users. The three types of SSH key algorithms that can be used are:
 - *ssh-ecdsa*

- ssh-ed25519
- ssh-rsa

**NOTE:**

- You can use both password and SSH keys to authenticate users.
- You can configure more than one SSH key to authenticate users.

5. Commit the changes and exit the configuration mode.

```
root@eop-primary# commit and-quit
warning: *** operating on 176.16.10.4 ***
warning: *** operation on 176.16.10.4 succeeds ***
warning: *** operating on 176.16.10.3 ***
warning: *** operation on 176.16.10.3 succeeds ***
warning: *** operating on 176.16.10.2 ***
warning: *** operation on 176.16.10.2 succeeds ***
commit complete
Exiting configuration mode
```

The user is successfully created.

6. Deploy the configuration on all the nodes in the cluster.

```
root@eop-primary> request paragon deploy user
Getting user-config configmap...
Added user: ['username']
Modified user: []
Deleted user: []
warning: *** operating on 176.16.10.4 ***
warning: *** operation on 176.16.10.4 succeeds ***
warning: *** operating on 176.16.10.3 ***
warning: *** operation on 176.16.10.3 succeeds ***
warning: *** operating on 176.16.10.2 ***
warning: *** operation on 176.16.10.2 succeeds ***
warning: *** operating on 176.16.10.1 ***
warning: *** operation on 176.16.10.1 succeeds ***
Deleting user-config configmap...
Creating new user-config configmap...
```

The configuration is deployed on all the nodes in the Routing Director cluster.

The system administrator must manually share the IP addresses and log in credentials with the user. The user can log in to Paragon Shell by using the IP addresses and login credentials to view and manage the Routing Director cluster based on the access privileges assigned to them.

7. (Optional) Confirm the user details.

```
root@eop-primary> configure
Entering configuration mode

[edit]
root@eop-primary# show system login
```

For every user created by using Paragon Shell, a Linux user with identical username, access privileges, and authentication method is created on every single node in the Routing Director cluster.



NOTE:

- When a commit is initiated from one node (node A), commit is first attempted on the rest of the nodes (node B, C, D) before it is attempted on the node on which the commit was initiated (node A). In case the commit fails on any of the nodes, the changes committed are rolled back on the nodes on which commit succeeded.
- Commit fails when multiple users make changes to the cluster configuration at the same time from different nodes. In such scenarios, we need to rollback all the committed changes on the other nodes and only commit changes from a single node to ensure that all the nodes in the cluster have the same configuration.

Modify User Information

System administrators and superuser can modify the access privileges of users and the authentication method used to authenticate the user.

To modify user details:

1. SSH to a node in the Routing Director cluster.
2. Log in to Paragon Shell.

If you are a system administrator, log in as the root user. If you are a superuser, log in with your login credentials shared by the system administrator.

You are placed in the operational mode.

3. Enter the configuration mode.

```
root@eop-primary> configure
Entering configuration mode

[edit]
```

4. Modify the access privilege for the user and the authentication method used to authenticate the user, as needed.

```
root@eop-primary# set system login user username authentication authentication-method
```

5. Commit the changes and exit the configuration mode.

```
root@eop-primary# commit and-quit
warning: *** operating on 176.16.10.4 ***
warning: *** operation on 176.16.10.4 succeeds ***
warning: *** operating on 176.16.10.3 ***
warning: *** operation on 176.16.10.3 succeeds ***
warning: *** operating on 176.16.10.2 ***
warning: *** operation on 176.16.10.2 succeeds ***
commit complete
Exiting configuration mode
```

The user is successfully created.

6. Deploy the configuration on all the nodes in the cluster.

```
root@eop-primary> request paragon deploy user
Getting user-config configmap...
Added user: []
Modified user: ['username']
Deleted user: []
warning: *** operating on 176.16.10.4 ***
warning: *** operation on 176.16.10.4 succeeds ***
warning: *** operating on 176.16.10.3 ***
warning: *** operation on 176.16.10.3 succeeds ***
warning: *** operating on 176.16.10.2 ***
warning: *** operation on 176.16.10.2 succeeds ***
warning: *** operating on 176.16.10.1 ***
warning: *** operation on 176.16.10.1 succeeds ***
```

```
Deleting user-config configmap...
Creating new user-config configmap...
```

The username of the user whose details are modified is displayed. The configuration is deployed on all the nodes in the Routing Director cluster.

Delete a User

To delete a user:

1. SSH to a node in the Routing Director cluster.
2. Log in to Paragon Shell.

If you are a system administrator, log in as the root user. If you are a superuser, log in with the login credentials shared by the system administrator.

You are placed in the operational mode.

3. Enter the configuration mode.

```
root@eop-primary> configure
Entering configuration mode

[edit]
```

4. To delete the user.

```
root@eop-primary# delete system login user username
```

5. Commit the changes and exit the configuration mode.

```
root@eop-primary# commit and-quit
warning: *** operating on 176.16.10.4 ***
warning: *** operation on 176.16.10.4 succeeds ***
warning: *** operating on 176.16.10.3 ***
warning: *** operation on 176.16.10.3 succeeds ***
warning: *** operating on 176.16.10.2 ***
warning: *** operation on 176.16.10.2 succeeds ***
commit complete
Exiting configuration mode
```

The user's access to Routing Director cluster is revoked.

6. Deploy the configuration on all the nodes in the cluster.

```

root@eop-primary> request paragon deploy user
Getting user-config configmap...
Added user: []
Modified user: []
Deleted user: ['username']
warning: *** operating on 176.16.10.4 ***
warning: *** operation on 176.16.10.4 succeeds ***
warning: *** operating on 176.16.10.3 ***
warning: *** operation on 176.16.10.3 succeeds ***
warning: *** operating on 176.16.10.2 ***
warning: *** operation on 176.16.10.2 succeeds ***
warning: *** operating on 176.16.10.1 ***
warning: *** operation on 176.16.10.1 succeeds ***
Deleting user-config configmap...
Creating new user-config configmap...

```

The user's access to Paragon Shell is removed. The configuration is deployed on all the nodes in the Routing Director cluster.

7. (Optional) Confirm that the user is logged out of the Routing Director cluster.

```

root@eop-primary> configure
Entering configuration mode

[edit]
root@eop-primary# show system login

```

Logging in to Paragon Shell as a New User

The system administrator manually shares the IP address and credentials (username and password) with the users to access Paragon Shell.

To log in to Paragon Shell as a new user:

1. SSH to a node in the Routing Director cluster.
2. Enter your username and password.

Based on your login class you are either placed in Paragon Shell or the Linux user shell.

The two login classes are:

- super-user—You are placed in the Paragon Shell.

- read-only—You are placed in the Linux user shell initially and prompted to re-enter your password. Re-enter your password to log in to Paragon Shell.

You are logged in to Paragon Shell.

Depending on your access privileges, you can view and manage Routing Director cluster configuration.

Retrieve User Information from a Recovered Node in the Routing Director Cluster

When a node is repaired and recovered in the Routing Director cluster, the user configuration on the node is lost. Users will not be able to use their user credentials to log in to the node and access the Linux user shell and Paragon Shell.



NOTE: Before you retrieve user configuration on the recovered node, ensure that pending commits are not present in any of the nodes in the Routing Director cluster, including in the recovered node.

To retrieve user configuration on the recovered node:

1. SSH to a node with the latest user configuration in the Routing Director cluster.

You are logged into Paragon Shell.

2. Enter the configuration mode.

```
root@eop-primary> configure
Entering configuration mode

[edit]
```

3. Execute the commit command to retrieve the latest cluster configuration.

```
[edit]
root@eop-primary# commit and-quit
commit complete
Exiting configuration mode
```

The user configuration is updated on all the nodes.

4. Deploy the user configuration file on all the nodes in the cluster.

```
root@eop-primary> request paragon deploy user recover true
Getting user-config configmap...
```

```

Added user: ['username']
Modified user: []
Deleted user: []
warning: *** operating on 176.16.10.4 ***
warning: *** operation on 176.16.10.4 succeeds ***
warning: *** operating on 176.16.10.3 ***
warning: *** operation on 176.16.10.3 succeeds ***
warning: *** operating on 176.16.10.2 ***
warning: *** operation on 176.16.10.2 succeeds ***
warning: *** operating on 176.16.10.1 ***
warning: *** operation on 176.16.10.1 succeeds ***
Deleting user-config configmap...
Creating new user-config configmap...

```

Users can now log in to the node using their existing Paragon Shell credentials. The access privileges and authentication method also remains the same.

Manage Your Routing Director Account

IN THIS SECTION

- [Change Account Information and Password | 126](#)
- [Enable Two-Factor Authentication | 126](#)
- [Enable E-Mail Notifications When SMTP is Configured | 127](#)
- [Delete Your Routing Director Account | 127](#)

You can manage your Routing Director account information from the My Account page. The My Account page displays the user's role in the organization as an immutable field. The user role is assigned by the superuser the superuser creates a user account in the organization.

To manage your account:

1. Log into Routing Director Web GUI.
2. Click the user account icon in the top-right corner of the GUI and choose **My Account** from the list.

The My Account page appears.

On the My Account page, you can see your account information—e-mail address, first name, last name, phone numbers, roles and authentication information. Network Admins can, in addition, view access control profiles assigned to them.

You can make changes as suggested in the following procedures:

- ["Change Account Information and Password" on page 126](#)
- ["Enable Two-Factor Authentication " on page 126](#)
- ["Enable E-Mail Notifications When SMTP is Configured" on page 127](#)
- ["Delete Your Routing Director Account" on page 127](#)

Change Account Information and Password

To change account information and password:

1. On the **My Account** page:

- a. In the Account Information section, change your e-mail address, name, and phone number, as necessary.
- b. In the **New Password** text field, type a new password to change the password.
The superuser configures the password policy for the organization. A password can contain up to 32 characters including special characters.

2. Click **Save**.

Routing Director updates your user account information and password.

Enable Two-Factor Authentication

To enable two-factor authentication:

1. On the **My Account** page, toggle the **Two Factor Authentication** button to enable two-factor authentication.

2. Click **Save**.

A message confirms updating your user data. A verify button appears near the two-factor authentication option.

3. Click **Verify**.

The Verification of Two Factor Authentication page appears displaying a QR code.

4. Open your authenticator application and click the add icon (+) to add a new account.

5. Scan the QR code displayed in Routing Director.

Your Routing Director account appears in your authenticator application.

6. Enter the token number from your authenticator application in the Verification of Two Factor Authentication page.

7. Click **Verify**.

A green check mark appears beside the Two Factor Authentication option on the My Account page. The two-factor authentication is active for your account. You can log out and log back in to Routing Director to verify if two-factor authentication is working.

Enable E-Mail Notifications When SMTP is Configured

If SMTP is configured in Routing Director, superuser, and network administrators can get e-mail notifications when alerts and alarms are generated for all or selected sites.

To enable e-mail notifications:

1. On the **My Account** page, click **Enable** in the Email Notification section.
The Enable Email Notifications page appears.
2. On the Enable Email Notifications page, do one of the following:
 - a. Click the **Enable Org Notifications** toggle button, to receive e-mail notifications for all sites in the organization.
 - b. Click the toggle button against a site to receive e-mail notifications specific to a site.
3. Click **Close**.

You are returned to the My Account page.

The Enable Email Notification section shows that you have enabled notifications for your current organization.

Delete Your Routing Director Account

To delete your account:

1. On the **My Account** page, click **Delete Account** present on the top-right corner of the page.
A confirmation message appears.
2. Click **Yes**.
Routing Director logs you out and deletes your account.



NOTE: After you delete your user account, Routing Director stores audit logs related to your actions for 30 days.

SEE ALSO

[Access Control Profile Overview](#) | 145

CHAPTER 8

Inventory Management

IN THIS CHAPTER

- [About the Inventory Page | 128](#)
- [Assign a Device to a Site | 133](#)
- [Release a Device | 133](#)
- [Replace a Device | 135](#)

About the Inventory Page

IN THIS SECTION

- [Tasks You Can Perform | 128](#)
- [Field Description | 130](#)

The Inventory page lists the devices in an organization grouped as routers, switches, and firewalls. You can view the device details such as hostname, model, device status, serial number, and so on, on this page.

To access the Inventory page, click **Inventory > Devices > Network Inventory** on the navigation menu.

Tasks You Can Perform

You can perform the following tasks on the Inventory page:

- View details of a device (router, switch, or firewall) present in the organization—To view details of a device, click the respective tab of the device, and click the **Details** icon that appears next to the check box beside a device name. The Device Details pane appears on the right side of the page displaying the basic device information and the site where the device is located. See [Table 40 on page 132](#).

- Adopt a device; see ["Add a Device to Routing Director" on page 188](#).
- Release a device; see ["Release a Device" on page 133](#).
- Export details of all the devices in a CSV format—To export details of all devices, on the respective tab, click the **Export** button. The details are exported to a CSV file that you can download to your local system.
- Assign one or more devices to a site; see ["Assign a Device to a Site" on page 133](#).
- Modify the operational state of a device—The device operational state is a label manually assigned to the device by the user. The label indicates at what predefined stage (onboarding, in service, and so on) the device is in the device life-cycle process. It does not indicate the device status.

To modify the operational state of a device after adopting the device:

1. Select a device and click **More > Edit Operational State** on the respective tab (Routers, Switches, or Firewalls).

The Edit Operational State page appears.

2. Select a value from the drop-down list and click **Ok**.

The value that you assigned for the operational state is visible in the Operational State column for that device.

- Assign a tag—A tag is a label in the *key:value* format that you assign to a device so that you can group the device based on its parameters. For example, you can assign a tag vendor:Juniper to all devices from Juniper Networks.

To assign one or more tags to a device:

1. Click **More > Manage Tags**.

The Manage Tags page appears.

2. In the **Tags** field, select a tag to be applied from the drop-down list that appears.

Alternatively, you can also create a new tag by clicking the **Create New Tag** link at the bottom of the page. If you are creating a new tag, ensure that:

- The key starts with a letter and is followed by letters, numbers, periods (.), hyphens (-), or underscores (_), and is between 1 and 200 characters long.
- The value includes alphanumeric characters, periods (.), hyphens (-), and underscores (_), and is between 1 and 200 characters long.

3. You can view all the tags applied to the device in the Tags column.

- Replace a device. See ["Replace a Device" on page 135](#).

- You can also perform the following tasks on this page:
 - Sort, resize, or re-arrange columns in a table (grid).
 - Show or hide columns in the table or reset page preferences, using the vertical ellipsis menu.
 - Search by using keywords—Click the search icon (magnifying glass), enter the search term in the text box, and press Enter. The search results are displayed on the same page.
 - Filter the data displayed in the table—Click the filter icon (funnel) and select whether you want to show or hide advanced filters. You can then add or remove filter criteria, save criteria as a filter, apply or clear filters, and so on. The filtered results are displayed on the same page.

For more information, see ["GUI Overview" on page 6](#).

- Set the number of rows to be displayed in the table—Set the number of rows (15, 30, 50, or 100) to be displayed in the Network Implementation Plan table in the **Display** drop-down list provided at the bottom right-corner of the table.

If there are multiple pages of rows, use the arrow heads to traverse through the pages or enter a page number to view the entries listed on that page.

Field Description

[Table 39 on page 130](#) lists the fields on the Inventory page.

Table 39: Fields on the Inventory Page (for the Routers, Switches, and Firewalls tabs)

Field	Description
ID	ID of the device in Routing Director.
Status	Status of the device: • Connected—The device is connected and assigned to a site in Routing Director. • Disconnected—The device is not connected or is connected but not assigned to a site in Routing Director.
Name	Name of the device. Click the name to view operational details (routing details, interface, management connectivity details, hardware issues, and so on) for the device.
IPv4 Address (for routers and firewalls)	Management IPv4 address assigned to the device.

Table 39: Fields on the Inventory Page (for the Routers, Switches, and Firewalls tabs) *(Continued)*

Field	Description
IPv6 Address (for routers and firewalls)	Management IPv6 address assigned to the device.
MAC Address (for switches)	MAC address assigned to the device.
Model	Device model; for example, ACX7100-48L, ACX7100-32C, and MX240.
Site	Site to which the device is assigned.
Serial Number	Serial number of the device.
Software Version	Version of the operating system (OS) installed on the device.
Product	Device type; for example, MX, ACX.
Vendor	Manufacturer of the device.
Operating System	OS installed on the device; for example, Junos and Junos Evolved.
Operational State	An optional state that can be assigned to the device. The predefined states are: • Undefined—The state of the device is not defined by Routing Director. This is the default state. • Onboarding—The device is in the process of being onboarded to Routing Director. • Ready for Service—The device is onboarded and ready to be moved to production. • In Service—The device is in production. • Maintenance—The device is undergoing maintenance. For example, software upgrade. • RMA—The device is getting replaced by another device.

Table 39: Fields on the Inventory Page (for the Routers, Switches, and Firewalls tabs) *(Continued)*

Field	Description
Tags	Tags assigned to the device. If there are more than one tag assigned to the device, an integer is present next to the displayed tag. For example, +2 in the Tags column indicates that there are two more tags in addition to what is currently listed under Tags. You can hover over the integer to view the additional tags assigned to the device.

Table 40: Fields on the *Device* Details Pane

Field	Description
General	
Name	Hostname of the device.
Model	Device model; for example, ACX7100-32C.
IP Address	Management IPv4 address assigned to the device.
Modified Time	Date and time when a device detail was modified.
Site Displayed only if a site is assigned to the device.	
Name	Name of the site where the device is installed.
Address	Address of the site where the device is installed.
Country Code	Country where the device is installed.
TimeZone	Timezone where the device is installed.

RELATED DOCUMENTATION

[About the Troubleshoot Devices Page](#) | 454

[Administration Workflow](#) | 55

Assign a Device to a Site

A site represents the location where the device is installed. Each device that is claimed (managed) by Routing Director must be assigned to a site for efficient management such as for applying policies.

To assign one or more devices to a site:

1. Navigate to **Inventory > Devices > Network Inventory**.
The Inventory page appears.
2. On the respective tab, select the device that you want to assign to a site and click **More > Assign to a Site**.
The Assign Devices to a Site page appears.
3. Select the site to assign the devices in the **Select Site** list and click **Done**.
The device is assigned to the selected site and the Site field on the Inventory page shows the site to which the device is assigned.

After the device is assigned to a site, you can apply all the device management functions on the device.

RELATED DOCUMENTATION

[Organization and Sites Overview](#) | 57

[About the Tags Page](#) | 141

Release a Device

SUMMARY

Use this procedure to release a device from the management for Routing Director.

A device should be released from the management of Routing Director if the device is no longer needed or when it is being decommissioned.

To release a device that is part of a network implementation plan, delete the device from the Network Implementation Plan and then use the **Release *device-type*** option present on the Inventory page (**Network Inventory > Inventory**).

To delete a device from a network implementation plan, edit the network implementation plan and delete the device from the Add Device tab of the Edit Network Implementation page (**Inventory > Device Onboarding > Network Implementation Plan**). See ["Edit a Network Implementation Plan" on page 260](#).

When you remove a device from the network implementation plan, all the configurations that were committed on the device through the plan are deleted, but the outbound SSH connection is retained.

Use the **Release *device-type*** option on the Inventory page to delete the outbound SSH connection from the device. Without the outbound SSH connection, the device is disconnected from Routing Director.

To release a device that is not part of a network implementation plan, use the **Release *device-type*** option. The **Release *device-type*** option deletes the outbound SSH command configured on the device disconnecting the device from Routing Director. Other configurations on the device are not affected.

To use the **Release *device-type*** option:

1. Navigate to **Inventory > Network Inventory**.

The Inventory page appears.

2. Select the device (under the appropriate tab) and click **Release *device-type***.

The Confirm *Device Release* page appears.

3. Do one of the following:

- If the device is not actively part of any services, click **Confirm** to release the device.
- If the device is actively a part of any service such as L3VPN, the list of such services is displayed. In such instances, you can either:
 - Manually remove the service configuration from the device and then release the device.
 - Click **Force Delete** to release the device forcefully.

The device is disconnected from Routing Director.

RELATED DOCUMENTATION

[Add a Device to Routing Director](#) | 188

Replace a Device

Sometimes, you might need to replace a device in your network (for example, during critical hardware failure). Routing Director provides the Replace Device option to help you replace a device without disrupting network services.

To replace a device, you must have another device of the same model installed in your network and onboarded to Routing Director. When you issue the command to replace a device, you are prompted to select the new device that would replace the old device.

Routing Director performs the following steps to replace a device:

1. Backs up configuration on the device you want to replace (old device).
2. Changes the device-specific details (for example, IP address and hostname) in the backed up configuration with the details of the new device.
3. Commits the backed up configuration from the old device on to the new device by using the CLI.
4. Sets the status of the original device to Undefined without offboarding the device.



NOTE: In this release, replacing a device is a beta feature. Routing Director can back up the configuration from the old device and commit the backed-up configuration on to the new device.

An audit log is generated when you replace a device and you can view it on the Audit Logs page (**Settings Menu > Audit Logs**).

Pre-requisites

The new device must be already onboarded to Routing Director and listed on the Inventory page. The device can be in the connected or disconnected state, but in the unassigned state (that is not assigned to any site).

To replace a device:

1. Click **Inventory > Devices > Network Inventory**.
The Inventory page appears.
2. Select the device that you want to replace and click **More > Replace Device**.
The Replace Router page appears.
3. On the Replace Router page, select a new router to replace the old router from the **Select New Router** drop-down list.
4. Click **Replace**.

The old device is replaced by the new device and the operational state of the old device is set to Undefined and the new device is set to In Service.

Audit Logs

IN THIS CHAPTER

- [Audit Logs Overview | 137](#)
- [About the Audit Logs Page | 138](#)

Audit Logs Overview

An audit log is a record of activities initiated by a user or by a process in a workflow that the user has initiated.

You can view a record of:

- User management events such as account creation, modification, deletion, and password change or reset.
User account lockout, unsuccessful authentication, and successful authentication are also logged.
- User-initiated activities such as creating, updating, or deleting a network implementation plan or a service and service resource from your network.

For example, a log is generated when a service order is created after you publish a service instance or when a network implementation plan is modified. You can view information about the order such as the type of order, timestamp, user who initiated the order, and so, in the audit log.

- System-run activities that are part of workflows in Routing Director. For example, committing the configurations defined in the network implementation plan on devices as part of the onboarding workflow by using the NETCONF protocol. Such tasks are recorded in the audit logs as system-initiated tasks even though the workflow is initiated by the user during the onboarding process.

Also, activities that are part of the provisioning workflows of L3VPN, L2VPN, and L2Circuit services are also recorded in the Audit Logs page. The audit log includes information such as time the workflow was initiated, workflow run ID, the user who created the service order, and so on.

Audit logs are useful in tracking and maintaining a history of user-initiated and system-run activities.

**NOTE:**

- Audit logging does not track device-initiated activities.
- Audit logs are cleared every 30 days.

Superusers and network administrators can view and filter audit logs to determine which users performed which actions at what time.

For example, a superuser or network administrator can use audit logs to see who:

- Added user accounts on a specific date.
- Accessed the organization and at what time.
- Updated or deleted an event (alert or alarm) template.
- Added or deleted a site.

RELATED DOCUMENTATION

[Administration Workflow | 55](#)

[Device Onboarding Workflow | 173](#)

[Service Provisioning Workflow | 648](#)

About the Audit Logs Page

IN THIS SECTION

- [Tasks You Can Perform | 139](#)
- [Field Descriptions | 139](#)

To access this page, select **Settings Menu > Audit Logs** on the banner. Superusers and network administrators can view and filter activities initiated by a user or by a process in a workflow that the user has initiated in an organization.

The Audit Logs page refreshes automatically and displays the latest logs.

Tasks You Can Perform

- View details of an audit log—Select an audit log and click **More > Detail** or click the **Details** icon on the left. The Details for Audit Log pane appears.

Hover over the **Period** drop-down list to filter the audit logs based on the time interval you select. You can choose Last 60 Minutes, Last 24 Hours, Last 7 Days, Today, Yesterday, This Week, or Custom (enter a custom time range).

- You can also perform the following tasks on this page:
 - Sort, resize, or re-arrange columns in a table (grid).
 - Show or hide columns in the table or reset page preferences, using the vertical ellipsis menu.
 - Filter the data displayed in the table—Click the filter icon (funnel) and select whether you want to show or hide advanced filters. You can then add or remove filter criteria, save criteria as a filter, apply or clear filters, and so on. The filtered results are displayed on the same page.

For more information, see ["GUI Overview" on page 6](#).

- Set the number of rows to be displayed in the table—Set the number of rows (15, 30, 50, or 100) to be displayed in the Network Implementation Plan table in the **Display** drop-down list provided at the bottom right-corner of the table.

If there are multiple pages of rows, use the arrow heads to traverse through the pages or enter a page number to view the entries listed on that page.

Field Descriptions

[Table 41 on page 139](#) describes the fields on the Audit Logs page.

Table 41: Fields on the Audit Logs Page

Field	Description
ID	Unique ID assigned to the log.
Timestamp	Date and time at which the audit log was recorded.
Username	Name and e-mail address of the user who initiated the task.
Source IP	IP address of the device from which the user initiated the task. For tasks that do not have an associated source IP address, this field is blank.

Table 41: Fields on the Audit Logs Page *(Continued)*

Field	Description
Message	Description of the logged task.
Event Type	Type of event for which the log is generated: • User Management: Includes user management events such as user creation and deletion, password reset, successful and unsuccessful user authentication, and so on. • Organization: Includes events related to an organization such as organization creation and deletion, user logging into the organization, and so on.
Site	Name of the site in which the task was initiated.
User Agent	Displays information about the browser the user used to access Routing Director GUI.
Job	Displays the Show job details link if a job is associated with the audit log activity. Click the link to search and display audit logs with the same Job ID.
Job ID	Unique ID assigned to the job.

RELATED DOCUMENTATION

[Manage Organization Settings](#) | 60

[About the Users Page](#) | 93

CHAPTER 10

Tags

IN THIS CHAPTER

- [About the Tags Page | 141](#)
- [Manage Tags | 143](#)

About the Tags Page

IN THIS SECTION

- [Tasks You Can Perform | 141](#)

A tag is a label in the *key:value* format assigned to a device or an interface (physical and logical). For example, you can assign a tag `interfacetype:Internet` to all the interfaces that connect to Internet, in a device.

A superuser or network administrator can view, add, edit, and delete tags.

You apply a tag so that you can group a set of devices and interfaces that have common features or requirements. For example, devices in the same city or a requirement to collect telemetry data. By using tags, you can group and filter devices and interfaces with specific tags.

To access this page, click **Inventory > Common Resources > Tags**.

Tasks You Can Perform

You can perform the following actions from the Tags page:

- View a list of all the tags configured in Routing Director. See [Table 42 on page 142](#).
- Add tags. See ["Add a Tag to Routing Director" on page 143](#).

- Edit tags. See ["Edit a Tag" on page 144](#).
- Delete tags. See ["Delete Tags" on page 144](#).
- Sort, resize, or re-arrange columns in a table (grid).
- Show or hide columns in the table or reset page preferences, using the vertical ellipsis menu.
- Search by using keywords—Click the search icon (magnifying glass), enter the search term in the text box, and press Enter. The search results are displayed on the same page.
- Filter the data displayed in the table—Click the filter icon (funnel) and select whether you want to show or hide advanced filters. You can then add or remove filter criteria, save criteria as a filter, apply or clear filters, and so on. The filtered results are displayed on the same page.
- Refresh the Tags table—Click the **Refresh** icon present at the bottom of the table to view the most recent updates made to the tags.
- Set the number of rows to be displayed in the table—Set the number of rows (15, 30, 50, or 100) to be displayed in the Network Implementation Plan table in the **Display** drop-down list provided at the bottom right-corner of the table.

If there are multiple pages of rows, use the arrow heads to traverse through the pages or enter a page number to view the entries listed on that page.

Table 42: Fields on the Tags Page

Field	Description
Key	A parameter associated with the device or interface. For example, interfacetype and service for interfaces and device-vendor for devices.
Value	Value for the key. • Internet or management for interfacetype • L3VPN, L2VPN, or broadband for service • Juniper or Cisco for device-vendor
Created Date	Date and time when the tag was created in the MM/DD/YYYY HH:MM format.
Created By	User who created the tag.

Table 42: Fields on the Tags Page *(Continued)*

Field	Description
Updated Date	Date and time when the tag was last updated in the MM/DD/YYYY HH:MM format.
Updated By	User who updated the tag.

Manage Tags

SUMMARY

Use this topic to add, edit, and delete tags that you can assign to devices. Tags are labels that you can assign to devices and interfaces (both physical and logical).

IN THIS SECTION

- [Add a Tag to Routing Director | 143](#)
- [Edit a Tag | 144](#)
- [Delete Tags | 144](#)

Add a Tag to Routing Director

To add a tag to Routing Director:

1. Click **Inventory > Common Resources > Tags**.

The Tags page appears.

2. Click the **+** (Add) icon to add a tag.

The Create Tag page appears.

3. In the **Tag** field, enter the tag in the *key:value* format. For example, `interfacetype:Internet`.

A *key* represents a property or a parameter of the device or interface. A key must start with a letter and can be followed by letters, numbers, periods (.), hyphens (-), or underscores (_), and between 1 to 200 characters long.

A *value* represents the value for that property or parameter. A value can include alphanumeric characters, periods (.), hyphens (-), and underscores (_) and between 1 to 200 characters long.

4. Click **OK**.

The tag is created and listed in the Tags table. The key is listed in the Key column and the value is listed in the Value column.

Edit a Tag

You can edit the key or the value or both of a tag. To edit a tag:

1. Click **Inventory > Common Resources > Tags**.

The Tags page appears.

2. Select the tag that you want to edit and click the **Edit** icon.

The Edit Tag page appears.

3. Edit the key, value, or both key and value of the tag.

A key must start with a letter and can be followed by letters, numbers, periods (.), hyphens (-), or underscores and between 1 and 200 characters long.

A value can include alphanumeric characters, periods (.), hyphens (-), and underscores and between 1 to 200 characters long.

4. Click **OK**.

The Edit Tag page closes and you can view the updates that you made to the tag in the Tags table.

Delete Tags

When you delete a tag, the list of resources (devices, network instances and network implementation plan) to which the tag was assigned is displayed.

To delete a tag:

1. Click **Inventory > Common Resources > Tags**.

The Tags page appears.

2. Select the tag and click **Delete** to delete the tag.

The Confirm Delete page appears.

3. Click **OK**.

The tag is removed from the Tags table. When you delete a tag that is assigned to a device or an interface, the tag is removed from the device or interface.

SEE ALSO

[About the Put Devices into Service Page | 264](#)

[About the Troubleshoot Devices Page | 454](#)

[About the Interfaces Tab | 466](#)

Attribute-Based Access Control

IN THIS CHAPTER

- [Access Control Profile Overview | 145](#)
- [About the Access Control Profiles Page | 151](#)
- [Add Access Control Profiles | 154](#)
- [Add Access Control Profile Policies | 155](#)
- [Edit and Delete Access Control Profiles | 165](#)

Access Control Profile Overview

SUMMARY


Attribute-based access control (ABAC) allows superusers to assign network administrators (users) access privileges granularly to Juniper Routing Director resources (for example, devices, network service, and tagging), and customize the access privilege assignment.

IN THIS SECTION

- [ABAC Workflow | 148](#)
- [Privilege Escalation | 150](#)
- [Benefits | 151](#)

With ABAC, superusers can grant permissions to network administrators (users in the rest of this topic) to perform a specific operation on a specific object. For example, superusers can assign a Network Admin an access control profile that allows the Network Admin (user) read and write access to core devices in a network, while restricting access to read-only for provider edge devices and vice-versa. Another example could be a user can access a Juniper device, but cannot access a Cisco device as shown in the [fig on page 146](#).

Figure 11: Restricting access using access control profile



In this release of Routing Director, a superuser can control access to the following resources—network services, devices, and tagging operations. If no profile is applied to a user or if a tag is not applied to a resource, then a user cannot access the resource.

Table 43 on page 146 lists the resources and the operations that can be access controlled in this release.

Table 43: Resources and operations planned for Access Control

Resource Name	Description	Operations				
		Create	Read	Update	Delete	Assign/Unassign
Device	Operations such as device onboarding, image upgrade, view device details, view device configuration , update and delete configuration , and so on	Onboarding a device is not access controlled	No	Yes	Yes	NA

Table 43: Resources and operations planned for Access Control (*Continued*)

Resource Name	Description	Operations				
Network Service	Configuring, updating, deploying and deleting Service instances,	No	No	Yes	Yes	NA
Tagging	Tag management and assignment	Yes	No	Yes	Yes	Yes

ABAC is implemented by using access control profiles. An access control profile comprises the following:

- **Policies**—Policies define the resources and conditions in a profile that determine whether a user is allowed or denied access to the resources. These conditions are based on tags assigned to the resources. Policies are defined in JSON format. Below is a sample of a policy *Residential Policy* that defines access condition for a *network service* that has a tag key `revenuegroup` and a tag value *Residential*.
- **Action**—The action (effect) whether to allow or deny access to resources when conditions in the policies are met.

When a user is assigned the following sample profile, `Sampleprofile`, the user is allowed write access (indicated by effect) to network services that have the tag `revenuegroup:Residential` assigned to them.

```
{
  "name": "Sampleprofile",
  "description": "Profile to allow access to l2vpn service design",
  "version": "1.0.0",
  "reference": "Non-mandatory way to input user's references like ticket number which
necessiated the change",
  "comment": "Non-mandatory way for user to input a comment like change reason",
  "effect": "Allow",
  "policies": [
    {
      "name": "Residential Policy",
      "apis": [
```

```

        "*"
      ],
      "resources": [
        "NetworkService"
      ],
      "conditions": {
        "StringResembles": {
          "revenuegroup": "Residential"
        }
      }
    }
  ]
}

```

By default, a superuser has the privilege to define tags, assign tags to resources, and also create, edit, and delete access control profiles. A superuser can provide the privilege to create, edit, and delete tags and assign tags to resources by assigning an appropriate profile to other users.

Routing Director logs the changes made to an access control profile and you can view the logs on the Audit Logs page (**Settings Menu > Audit Logs**).


ABAC Workflow

The workflow to allow a user access to resources using ABAC is as follows:

1. Create an access control profile. See ["Add Access Control Profiles" on page 154](#).
2. Assign the profile to a user. See ["Invite Users" on page 109](#).
3. Enforce the profile while handling requests from the user. Routing Director allows or denies access to resources based on the action (effect) defined in the profile.

A user can be assigned more than one access profile containing multiple policies. [Figure on page 149](#) shows how the behavior when more than one access control profiles and policies are assigned to a user.

Figure 12: Behavior when Multiple Profiles and Policies are Assigned



When multiple profiles are assigned, the user is allowed or denied access to a resource as listed in [table on page 146](#):

Table 44: Access when multiple profiles are assigned to a user

Access Control Profile Effect	Routing Director Action
All profiles have Allow effect.	If a resource meets the conditions in any one of the profiles, then a user is allowed access to resources defined in the profile.

Table 44: Access when multiple profiles are assigned to a user *(Continued)*

Access Control Profile Effect	Routing Director Action
Some profiles have Allow and some have Deny effect	Routing director first evaluates the profiles with "Deny" effect and if any of them evaluates to True, access is denied to the user. If conditions in none of the profiles with "Deny" effect is satisfied, then the profiles with "Allow" effect are evaluated. If any one of those profile evaluates to True, then the user is granted access to the resource.



NOTE: You cannot have all profiles with Deny effect.

Privilege Escalation

A profile that allows tagging a resource with one or more keys, when combined with other profiles that allow access to a resource tagged with those same keys grants users greater privileges. This is because the user can tag any resource with those keys and gain access to it.

For example, a profile P1, with condition `region:US`, would not allow a user to access devices tagged `region:canada`. But a profile P2, that allows tagging devices with the `region` key, in combination with profile P1, would allow the user to tag devices with the `region:US` tag or change the tag value from `canada` to `US` and then gain access to them.

To warn that a profile includes a tagging policy, Routing Director:

- Displays privilege level as **High** on the Access Control Profiles page (**Settings Menu > Control Profiles**).
- Displays a Warning sign next to the profile in the **Access Control Profile** drop-down list (for example, while assigning a profile to a user on the Add Users page).
- Displays a Warning message **Profile with high privilege level** when the profile is selected.

To further limit the capability of users assigned a tagging profile, a reserved key `internal.tag.constraints` is supported. The reserved key restricts access to resources, ensuring that a user would only tag resources tagged with those specified by `internal.tag.constraints` and not be able to tag any arbitrary resource.

For example, the following condition in a policy would allow a user to access any resource tagged with department:IT, CT0, or CF0 AND vendor:Juniper, and tag such a resource only with keys region and type.

```
"conditions": {
  "StringEquals": {
    "internal.tag.constraints": {
      "department": ["IT", "CT0", "CF0"],
      "vendor": ["Juniper"]
    },
    "internal.tag.keys": ["region", "type"]
  }
}
```

Benefits

- **Enhanced Granular Control:** ABAC (Attribute-Based Access Control) allows a superuser to implement more detailed and specific controls over resources, ensuring that access permissions are tailored to individual attributes and contexts.
- **Streamlined Administrative Experience:** By using ABAC, the process of managing access for new users and resources is simplified and made efficient.
- **Improved Tag Security and Reliability:** Controlling the permission to tag resources ensures that tags remain secure and reliable. This is crucial for maintaining the integrity of Routing Director use-cases and access control mechanisms within the system.

About the Access Control Profiles Page

SUMMARY

IN THIS SECTION

- [Tasks You Can Perform | 152](#)
- [Field Description | 152](#)

To access the Access Control Profiles page, click **System Settings > Control Profile** on the banner.

Tasks You Can Perform

You can perform the following tasks on the Access Control Profiles page:

- View access control profile details—To view details of an access control profile, select the access control profile and click **More > Detail**. Alternatively, hover over the profile name and click the **Details** icon that appears.

The Details for Profile: *profile-name* pane appears on the right side of the Access Control Profiles page displaying the details of the profile.

See [Table 45 on page 153](#) for information about the Access Control Profiles parameters.



NOTE: A Super User can view details of all the access control profiles, while a Network Admin can view the access control profiles assigned to the Network Admin.

- Add access control profiles. See ["Add Access Control Profiles" on page 154](#).
- Edit and delete access control profiles. See ["Edit and Delete Access Control Profiles" on page 165](#).
- You can also perform the following tasks on this page:
 - Sort, resize, or re-arrange columns in a table (grid).
 - Show or hide columns in the table or reset page preferences, using the vertical ellipsis menu.
 - Search by using keywords—Click the search icon (magnifying glass), enter the search term in the text box, and press Enter. The search results are displayed on the same page.
 - Filter the data displayed in the table—Click the filter icon (funnel) and select whether you want to show or hide advanced filters. You can then add or remove filter criteria, save criteria as a filter, apply or clear filters, and so on. The filtered results are displayed on the same page.

For more information, see ["GUI Overview" on page 6](#).

- Set the number of rows to be displayed in the table—Set the number of rows (15, 30, 50, or 100) to be displayed in the Access Control Profiles table in the **Display** drop-down list provided at the bottom right-corner of the table.

If there are multiple pages of rows, use the arrow heads to traverse through the pages or enter a page number to view the entries listed on that page.

Field Description

["Table 1" on page 153](#) lists the access control profile parameters.

Table 45: Access Control Profile Parameters

Field	Description
Name	Name of the access control profile.
Description	Description of the profile.
Reference	Reference associated with the access control profile. For example, a ticket number for which the access control profile is created or modified.
Comment	Comment entered for the access control profile; for example, while editing the profile.
Effect	Displays, whether the profile would allow or deny access to resources defined in the profile. • Allow—When assigned to a user, the profile allows access to resources defined in its policies. • Deny—When assigned to a user, the profile denies access to resources defined in its policies.
Privilege Level	Indicates the privilege level of the profile: • High—The profile includes one or more tagging policies. When assigned to a user, the user can tag resources defined in the profile policies. • Low—The profile does not include a tagging policy.
Policies	Policies included in the profile. You can view the details of the policy in JSON format in the Detail of Profile: <i>profile-name</i> pane. In the Details pane, click the arrow or the ellipses under Details for: <i>policy name</i> to see the resources and conditions defined in the policy.

Add Access Control Profiles

SUMMARY

Use this topic to add access control profiles to Routing Director.

Before you add access control profiles, you should have assigned tags to resources that you want to restrict access to.

To add an access control profile:

1. Click **Settings Menu > Control Profiles**.

The Access Control Profiles page appears.

2. Click the + (Add) icon.

The Add Control Profile page appears.

3. Enter values referring to table.

4. Click **Save**.

A message indicating that the control profile is created appears and the control profile is listed on the Access Control Profiles page.

[Table on page 154](#) lists parameters on the Add Control Profiles page.

Table 46: Access Control Profile Parameters

Field	Description
Name	Enter a name for the access control profile. A profile name can include any type of character but must not exceed the 64-character limit.
Description	Enter a description for the control profile.
Version	The version of control profile. This field is read-only in this release.
Reference	Enter a reference associated with the access control profile. For example, a ticket number for which the access control profile is being created or modified.

Table 46: Access Control Profile Parameters *(Continued)*

Field	Description
Comment	Enter a comment for the access control profile; for example, a comment entered when editing the profile.
Effect	Select the effect of the profile. Possible values: • Allow—The profile, when applied to a user, allows the user access to resources defined in the profile policies. • Deny—The profile, when applied to a user, denies the user access to resources defined in the profile policies.
Policies	The policies included in the access control profile. Enter the policy in the JSON format. See "Add Access Control Profile Policies " on page 155 for details.

What's Next

Apply the control profile to network administrators to provide granular access to Routing Director resources.

Add Access Control Profile Policies

SUMMARY

This topic describes the operators, conditions, and rules for adding an access control policy, with examples.

IN THIS SECTION

- [Condition Operators Supported in a Policy | 156](#)
- [Restrict Tagging Privilege using Condition Overrides | 158](#)
- [Rules for Adding Policies to a Profile | 161](#)

An access control policy in an access control profile defines the resources and the conditions for accessing the resources. You can define conditions for accessing the following resources in an access control policy:

- Devices

- Network services
- Tagging operations

 **NOTE:** The resources must be used as Device, NetworkService, and Tagging in a policy.

For details of operations that can be access controlled on devices, network services, and tagging operations, see [Table 43 on page 146](#).

For information on adding an access control profile, see ["Access Control Profile Overview" on page 145](#).

Condition Operators Supported in a Policy

[Table 47 on page 156](#) lists the operators that you can use in an access control policy.

Table 47: Access Control Policy Operators

Operator	Description	Example
StringEquals	Use for string equals comparison of tags assigned to resources. More than one tag in a condition results in AND operation of the tags. When the example condition is used in a policy and assigned to a network administrator (user in rest of the topic), the user will have access to all devices and network services with tag device:Juniper and site:NY.	<pre>... "conditions": { "StringEquals": { "device": "Juniper", "site": "NY" } } ...</pre>

Table 47: Access Control Policy Operators *(Continued)*

Operator	Description	Example
ForAnyValues:StringEquals	Use for string equals comparison of tags assigned to resources. More than one tag in a condition results in OR operation of the tags. When the example condition is used in a policy and assigned to a user, the user will have access to devices and services tagged with device:Juniper or site:NY.	<pre>... "conditions": { "ForAnyValues:StringEquals": { "device": "Juniper", "site": "NY" } } ...</pre>
StringResembles	Use for string resembles comparison of tags assigned to resources. You can use the wildcard character asterisk (*) at the beginning or at the end of the tag key and tag value. More than one tag in a condition results in AND operation of the tags. When the example condition is used in a policy and assigned to a user, the user will have access to devices and services tagged with keys starting with depart and values ending with st, such as department:Test.	<pre>... "conditions": { "StringResembles": { "depart*": "*st" } } ...</pre>

Table 47: Access Control Policy Operators (*Continued*)

Operator	Description	Example
ForAnyValues:StringResembles	Use for string resembles comparison of tags assigned to resources. You can use the wildcard character (*) at the beginning and end of the tag key and tag value. More than one tag in a condition results in OR operation of the tags. When the example condition is used in a policy and assigned to a user, the user will have access to devices and services tagged with keys starting with depart and values ending with st, such as department:Test or state:in-maintenance or state:under-maintenance.	<pre>... "conditions": { "ForAnyValues:StringResembles": { "depart*": "*st", "state": "*-maintenance", } } ...</pre>



NOTE: StringResembles and ForAnyValues:StringResembles operator supports wildcard character (*) only at the beginning or at the end of a tag key or tag value. If you use * as the tag value in a condition, then any value is matched. For example:

- us-* for matching us-west, us-east, and so on
- *engineering for matching test-engineering and support-engineering, and so on)
- * for matching any value like us, europe, or K2 peak

Restrict Tagging Privilege using Condition Overrides

Routing Director allows a superuser to enable users to assign tags. See ["Privilege Escalation" on page 150](#) for details.

To restrict the privilege of assigning tags, Routing Director provides key override constructs. Use the key override constructs as follows:

- `internal.tag.keys`—Use to provide a list of tag keys as conditions when tagging is specified as a resource in a policy.

The below sample policy, when used in a profile and assigned to a user, will allow the user to :

- Assign tags with `department` and `role` keys to resources. For example, `department:Sales` and `role:Sales`
- Create tags with key `department` and `role`.

```
"policies": [
  {
    "name": "policy1",
    "apis": [
      "*"
    ],
    "resources": [
      "Tagging"
    ],
    "conditions": {
      "StringEquals": {
        "internal.tag.keys": [
          "department",
          "role"
        ]
      }
    }
  }
]
```

- `internal.tag.constraints`—Use to restrict the tagging operation to a specific set of resources instead of allowing unrestricted tagging on all resources.

For example, the following condition in a policy would allow a user to access any resource tagged with `department:IT`, `CTO`, or `CF0` AND `vendor:Juniper`, and tag such a resource only if the resource had keys `region` or `type`.

```
"conditions": {
  "StringEquals": {
    "internal.tag.constraints": {
      "department": ["IT", "CTO", "CF0"],
      "vendor": ["Juniper"]
    }
  }
}
```



```

    },
    "internal.tag.keys": ["region", "type"]
  }
}

```



NOTE: Keys specified in `internal.tag.keys` cannot be the same as any of the keys specified in `internal.tag.constraints`. This is to prevent privilege escalation. See ["Privilege Escalation" on page 150](#) for details on privilege escalation.

- `internal.network-service.design-id`—Use to provide service design IDs such as infrastructure, elan-evpn-csm, l3vpn, and so on.

For example, when you assign the following profile to a user, the user can create, update, or delete L3VPN service instances.

```

{
  "name": "profile1",
  "description": "Profile to allow access to l3vpn service instances",
  "version": "1.0.0",
  "reference": "Non-mandatory way to input user's references like ticket number which
necessitated the change",
  "comment": "Non-mandatory way for user to input a comment like change reason",
  "effect": "Allow",
  "policies": [
    {
      "name": "policy1",
      "apis": [
        "*"
      ],
      "resources": [
        "NetworkService"
      ],
      "conditions": {
        "StringEquals": {
          "internal.network-service.design-id": [
            "l3vpn"
          ]
        }
      }
    }
  ]
}

```

```
]
}
```

Rules for Adding Policies to a Profile

You must consider the following points while creating an access control policy. Otherwise, Routing Director considers the policies as ambiguous:

- An access control policy with Tagging resource must contain at least one `internal.tag.keys` construct. The `internal.tag.keys` construct should not be combined with other tagging conditions in the same condition block. This is to prevent privilege escalation. If you combine `internal.tag.keys` construct and other tagging conditions in the same condition block, Routing Director considers the policy as ambiguous and does not process it.

For example, when the following profile is applied to a user, the user can assign department and vendor tags to any device or network service. However, the `internal-tag-constraints` restricts the user to assign the department and vendor tags to resources with `region:US`, `region:Canada` or `region:France` tag.

```
{
  "name": "profile1",
  "description": "Profile to allow tagging resources with department and vendor tags.",
  "version": "1.0.0",
  "reference": "Non-mandatory way to input user's references like ticket number which
necessitated the change",
  "comment": "Non-mandatory way for user to input a comment like change reason",
  "effect": "Allow",
  "policies": [
    {
      "name": "policy1",
      "apis": [
        "*"
      ],
      "resources": [
        "Tagging"
      ],
      "conditions": {
        "StringEquals": {
          "internal.tag.keys": [
            "department",
            "vendor"
          ],
          "internal.tag.constraints": {
```

```

        "region": ["US", "Canada", "France"],
      }
    }
  }
]
}

```

- If the resources list includes a network service, then the condition must include `internal.network-service.design-id` to indicate the service design ID that a user can access.

For example, when the policy is applied to a user, the user can create, update, or delete a service instance of service design type *elan-evpn-csm* tagged with the `region:europa` tag.

```

"policies": [
  {
    "name": "policy1",
    "apis": [
      "*"
    ],
    "resources": [
      "NetworkService"
    ],
    "conditions": {
      "StringEquals": {
        "internal.network-service.design-id": [
          "elan-evpn-csm"
        ],
        "region": "europa"
      }
    }
  }
]

```

- If the resources list includes a device, the policy conditions should include at least one tag for the device.

For example, when the policy is applied to a user, the user can access devices with tag `region:Europe` or `MiddleEast` AND `role:PE`.

For example:

```
"policies": [
  {
    "name": "policy1",
    "apis": [
      "*"
    ],
    "resources": [
      "Device"
    ],
    "conditions": {
      "StringEquals": {
        "region": [
          "Europe",
          "MiddleEast"
        ],
        "role": "PE"
      }
    }
  }
]
```

- Use of wildcards for `internal.network-service.design-id` construct when using `StringResembles` and `ForAnyValues:StringResembles` operator is not supported.

For example, Routing Director would not be able to resolve the following policy:

```
"policies": [
  {
    "name": "policy1",
    "apis": [
      "*"
    ],
    "resources": [
      "NetworkService"
    ],
    "conditions": {
      "StringResembles": {
        "internal.network-service.design-id": [
          "el*"
        ]
      }
    }
  }
]
```

```

        ],
        "region": "europe"
    }
}
]

```

- Do not start a tagging policy condition with the `ForAnyValues` operator.

For example, Routing Director would not be able to resolve the following policy:

```

"policies": [
  {
    "name": "policy1",
    "apis": [
      "*"
    ],
    "resources": [
      "Tagging"
    ],
    "conditions": {
      "ForAnyValues:StringEquals": {
        "department": [
          "Test"
        ]
        "instance": [
          "POC"
        ]
      }
    }
  }
]

```

- Use wildcard comparison only with **StringResembles** or **ForAnyValues:StringResembles** operators. With **StringEquals** or **ForAnyValues:StringEquals**, the comparison is always string comparison even if you use wildcard characters.
- The same tag keys cannot be used in the `internal.tag.keys` and `internal.tag.constraints`. This is to prevent privilege escalation. Use `internal.tag.keys` only with the `StringEquals` operator.

For example, the following condition in a policy would allow a user to access any resource tagged with `department:IT`, `CTO`, or `CF0` AND `vendor:Juniper`, and tag such a resource only with keys `region` and `type`.

```
"conditions": {
  "StringEquals": {
    "internal.tag.constraints": {
      "department": ["IT", "CTO", "CF0"],
      "vendor": ["Juniper"]
    },
    "internal.tag.keys": ["region", "type"]
  }
}
```

Edit and Delete Access Control Profiles

SUMMARY

Use this topic to edit or delete an access control profile.

IN THIS SECTION

- [Edit an Access Control Profile | 165](#)
- [Delete an Access Control Profile | 166](#)

Edit an Access Control Profile

You must be a superuser to edit an access control policy. All fields except **Version** can be edited.

When you edit an access control profile that is assigned to one or more users, the changes take effect immediately on the users.

To edit an access control policy:

1. Click **Settings Menu > Control Profiles** on the banner.
The Access Control Profiles page appears.
2. Click the profile that you want to edit.
The Edit Access Control Profile page appears.
3. Edit the fields by referring to [Table 46 on page 154](#).
4. Click **Save**.

You are returned to the Access Control Profiles page, where you can verify the changes in the Access Control Profiles table.

Delete an Access Control Profile

You must be a superuser to delete an access control profile. An access control profile assigned to a user cannot be deleted. Therefore, before you delete an access control profile, ensure that the profile is not assigned to any user.

To edit an access control policy:

1. Click **Settings Menu > Control Profiles** on the banner.

The Access Control Profiles page appears.

2. Click the profile that you want to delete.

A confirmation message appears.

3. Click Yes.

The profile is deleted, and the profile is no longer listed on the Access Control Profiles page.

3

PART

Device Life Cycle Management

- Introduction | **168**
 - Day-Wise Activities for Device Life Cycle Management | **177**
 - Field Technician User Interface | **199**
 - Onboarding Profiles | **202**
 - Plan Device Onboarding | **230**
 - View Device Onboarding | **264**
 - Device Management | **355**
-

CHAPTER 12

Introduction

IN THIS CHAPTER

- [Device Life-Cycle Management Overview | 168](#)
- [Device Onboarding Overview | 171](#)
- [Device Onboarding Workflow | 173](#)

Device Life-Cycle Management Overview

IN THIS SECTION

- [Onboard a Device | 169](#)
- [Manage and Monitor a Device | 169](#)
- [Decommission a Device | 170](#)
- [Benefits of Device Life Cycle Management | 171](#)

Device life-cycle management (LCM) in Routing Director is divided into various tasks that you perform as Day -2, Day -1, Day 0, Day 1, and Day 2 activities. The tasks are divided so that you follow a structured process to onboard, manage, and offboard devices. The activities for managing a device life cycle are divided as:

- Day -2 activities in which a network architect plans the device role and device configuration for that device role. See ["Add Network Resource Pools and Profiles \(Day -2 Activities\)" on page 177](#).
- Day -1 activities in which a network planner prepares a plan for onboarding the device to Routing Director. See ["Prepare for Device Onboarding \(Day -1 Activities\)" on page 178](#).
- Day 0 activities in which a field technician installs the device. The field technician or a network administrator gets Routing Director to manage the device. See ["Install and Onboard the Device \(Day 0 Activities\)" on page 179](#).

- Day 1 and Day 2 activities in which a network administrator monitors the health and functioning of the device and moves the device to production. See ["Move a Device to Production \(Day 1 and Day 2 Activities\)" on page 197](#).

Onboard a Device

You can use Routing Director to onboard:

- New devices that you procure for your network (greenfield devices).

You onboard greenfield devices by using a network implementation plan, which includes the management configurations (IP address, hostname, and so on) and infrastructure configurations (routing protocol configurations). Routing Director applies the following configurations on a device by using a network implementation plan:

- Basic device-level configurations (for example, IP address configurations, hostname, and software image to be used) and routing protocols (for example, IS-IS, OSPF, BGP, RSVP, LDP, and PCEP).
- Configuration for links with neighboring devices. The neighboring devices are devices that are a part of the same network implementation plan.
- Configuration for performing health checks, connectivity checks, and running trust scans.

These configurations are applied when you onboard a device to Routing Director either through the field tech UI or by committing outbound SSH commands on the device.

- Devices that exist in your network (brownfield devices).

You onboard brownfield devices by committing outbound SSH commands for connecting with Routing Director, on the device. Routing Director provides the SSH commands that you can copy and commit on the device. The onboarding of a device by committing the outbound SSH commands is referred to as adopting a device.

See ["Device Onboarding Overview" on page 171](#).

Manage and Monitor a Device

After you onboard a device, you can manage a device's inventory, apply licenses, backup and restore device configurations, upgrade software, reboot the device, and access the device CLI. See ["Device Onboarding Workflow" on page 173](#).

While Routing Director provides automated solutions for managing configurations, device monitoring, and periodic Trust scans for greenfield devices, Routing Director also provides the conventional device LCM solutions for brownfield devices.

For a greenfield device, to upgrade software, you update the software version to be applied on the device in the device profile or the network implementation plan used to onboard the device. Similarly, links and basic configurations that were committed on a device by using the network implementation plan can be updated by editing the network implementation plan and profiles used to onboard the device. You can also use configuration templates to apply advanced configurations on the device.



NOTE: After onboarding, you can upgrade the software on a greenfield device by using the Software Upgrade option (**Observability > Health > Troubleshoot Devices > More > Upgrade**) as well, without having to update the network implementation plan.

In addition, Routing Director executes a predetermined set of tasks (based on the configurations in the plan and profiles) for automatic monitoring and operations of the greenfield devices right from when the device is onboarded. For example, when you enable BGP or RSVP protocols in the profiles, Routing Director monitors the functioning of the BGP and RSVP protocols and displays any alerts or alarms related to the functioning of the protocols, right after the device is onboarded.

Routing Director GUI provides an integrated view of all the information about a device. On the Device-Name page (**Inventory > Onboarding Dashboard > *Device-Hostname***), you can view general details, connectivity details, results of trust scans, and key performance indicators (KPIs), and assess the functioning of the device. You can also upgrade software and perform a backup of the device configurations from the same page.

For brownfield devices, Routing Director provides options for software upgrade, adding licenses, applying configurations by using configuration templates, and backing up configurations under the **Observability > Health > Troubleshoot Devices** menu.

Decommission a Device

When you want to decommission (offboard) a greenfield device, you can:

- Use the network implementation plan that you are using to manage a device to decommission the device. See ["Offboard a Network Implementation Plan" on page 259](#).

When you use a network implementation plan to offboard, device configurations are deleted, but the outbound SSH configuration is retained. You must delete the outbound SSH configuration for Routing Director to disconnect from the device. See ["Release a Device" on page 133](#).

- Use the Release option to delete the outbound SSH configuration so that Routing Director disconnects from the device, See ["Release a Device" on page 133](#).

In this case, the other configurations committed on the device are retained. You must access the device CLI and manually delete the configurations.

To decommission a brownfield device, you simply use the Release option in Routing Director to delete the outbound SSH configuration on the device. See ["Release a Device" on page 133](#).

Benefits of Device Life Cycle Management

- Provides an automated solution for managing the life cycle of new devices procured for a network.
- The profiles and network implementation plan that are used to onboard and manage multiple devices considerably reduce the time and effort taken to onboard and manage devices. For example, to upgrade software running on five devices, you can simply edit the software version in the network implementation plan used for onboarding the devices and publish the plan. Routing Director updates the software on the devices to the version you mention in the plan.

RELATED DOCUMENTATION

| [Device Management Workflow](#) | 355

Device Onboarding Overview

Device onboarding refers to the steps that you must perform to enable Routing Director to manage the devices in your network. Device onboarding involves different personas in an organization performing different tasks to onboard devices.

A network architect prepares to add devices to the network and decides the roles for each device in the network. Based on the device role, the network architect creates resource pools, device profiles, and interface profiles.

Resource pools include values for network resources [IP addresses, loopback addresses, BGP cluster IDs, segment identifiers (SIDs), autonomous system number, and so on] that Routing Director can assign to the devices when automatic configuration is specified for the resources. See the ["Add Network Resource Pools" on page 676](#) for more details.

The device profiles include configurations such as IP loopback address, router ID, the software image to be used, and some routing protocols (such as BGP). The interface profiles include the routing protocol (IS-IS, OSPF, RSVP, and LDP) configurations. The network architect can also specify compliance and connectivity checks to be performed during device onboarding. See ["Device and Interface Profiles Overview" on page 202](#) for more details.

A network planner uses these profiles to create a plan (referred to as network implementation plan) for onboarding devices. In the plan, the network planner assigns the device and interface profiles to the

devices to be onboarded. The planner can also configure links between the devices included in the plan. See ["Network Implementation Plan Overview" on page 230](#) for more details.

The planner also adds information about the type of pluggables and cables to be used for each port on a device.

Routing Director commits configurations defined in the device and interface profiles, and the network implementation plan on the device during device onboarding. You can use the profiles and plan to also add configurations after a device is onboarded. For example, if a plan has an RSVP LSP configured from a device to all the provider edge (PE) devices, an LSP is configured from the device to all the PE devices that are present in the network during onboarding and also, to any PE device that might be added to the network after the device is onboarded.

After a device is onboarded and brought to production, you can use the network implementation plan to manage the devices. For example, if you want to upgrade software on all the devices in the plan, you specify the software version to be installed in the plan and push the updates on to the devices (known as publish). Routing Director updates the software that is installed on the devices to the version you specified in the plan.

Apart from the field technician, the superuser or network administrator can also onboard devices (both greenfield and brownfield devices) by committing the outbound SSH commands that Routing Director provides.

After the device connects with Routing Director, you can manage the device by using Routing Director. See ["Device Management Workflow" on page 355](#).

Benefits

- Routing Director facilitates faster deployment of devices to the network by committing device configurations and checking the health and connectivity of the devices during onboarding.
- The network implementation plan provides an easy way to upgrade software or modify configurations on multiple devices at the same time.

RELATED DOCUMENTATION

[Add a Device Profile | 208](#)

[Add an Interface Profile | 222](#)

[Add a Network Implementation Plan | 241](#)

Device Onboarding Workflow

The workflow for onboarding a new device (greenfield device) includes creating network resource pools, device and interface profiles, and a network implementation plan. The network implementation plan includes instructions about the type of pluggables and cables that a field technician must use for the device ports.

[Table 48 on page 173](#) lists the different personas and the roles in Routing Director that are involved in onboarding a device.

Table 48: Persona and Roles Involved in Device Onboarding

Persona	Role in Routing Director
Network architect	Super User or Network Admin
Network planner	Super User or Network Admin
Field Technician	Installer
NOC Engineer (Network administrator)	Super User or Network Admin

Onboard a New Device

To onboard a new device to Routing Director:

1. A network architect creates network resource pools for automatic assignment of values to the resource pools (IP addresses, segment identifiers, BGP cluster IDs, and so on). See ["Add Network Resource Pools" on page 676](#).
2. The network architect decides the configurations that must be committed on the devices to be onboarded and creates the following profiles:
 - Device profiles. See ["Add a Device Profile" on page 208](#).
 - Interface profile. See ["Add an Interface Profile" on page 222](#).

The network architect can add device and interface profiles to suit specific needs; that is, create profiles with configurations that can be committed to all the devices or selected devices in a network.

3. A network planner creates a network implementation plan for onboarding the device. See ["Add a Network Implementation Plan" on page 241](#).
4. At the site, the field technician unpacks the device and mounts it on a rack.

For instructions on how to mount a device, see the corresponding device Hardware Guide or the [Quick Start Guide](#) in the [Documentation](#) site. To access the Hardware Guide or the Quick Start Guide of a device, on the homepage of the Documentation site, under **Products by Category**, click **View More > *Device-Model*** in the **Routing** section.

5. The device can be onboarded to Routing Director in one of the following ways:
 - A superuser or network administrator can adopt the device; see ["Add a Device to Routing Director"](#) on page 188.
 - The field technician can use the field technician UI to onboard the device. See ["Working with Field Technician UI Pages"](#) on page 200.




NOTE: Use of the field technician UI to onboard the device is an experimental feature.

6. A network administrator applies additional configurations on the device for production by using configuration templates. See ["Deploy a Configuration Template to a Device"](#) on page 390.
7. The network administrator tests and certifies that the device is ready for production.
8. The network administrator moves the device to production. See ["Move a Device to Production"](#) on page 268.
9. The network administrator monitors the functioning of the device from the Routing Director GUI. See ["View Results of Automated Device Tests"](#) on page 268.

[Figure 13 on page 175](#) shows the device onboarding workflow in Routing Director for a new device (greenfield).

Figure 13: Device Onboarding Workflow



Adopt a Brownfield Device

You can onboard a brownfield device to Routing Director without adding the device to a network implementation plan. In this case, you will be able to onboard the device to Routing Director, but cannot view the device and health metrics for the device. However, you can add the device to a network implementation plan and then adopt the device, in which case, you can monitor the device and interface health of the device.

To onboard a brownfield device by adding it to a network implementation plan:

1. Create a device and interface profile with Trust enabled (in device profile) and all other configurations disabled. See ["Add a Device Profile" on page 208](#) and ["Add an Interface Profile" on page 222](#).
2. Create a network implementation plan and add the device profile and interface profile to the network implementation plan. See ["Add a Network Implementation Plan" on page 241](#).
3. Add the device to the network implementation plan.
4. Publish the network implementation plan. See ["Publish a Network Implementation Plan" on page 258](#).

After a device is onboarded to Routing Director and listed in the Inventory page (**Inventory > Network Inventory**), you can monitor the device and interface health from the Device accordion and Interface accordion on the *device-name* page. See ["About the Device-Name Page" on page 461](#).

RELATED DOCUMENTATION

[Device and Interface Profiles Overview | 202](#)

[Device Onboarding Overview | 171](#)

[Add Network Resource Pools and Profiles \(Day -2 Activities\) | 177](#)

[Prepare for Device Onboarding \(Day -1 Activities\) | 178](#)

Day-Wise Activities for Device Life Cycle Management

IN THIS CHAPTER

- [Add Network Resource Pools and Profiles \(Day -2 Activities\) | 177](#)
- [Prepare for Device Onboarding \(Day -1 Activities\) | 178](#)
- [Install and Onboard the Device \(Day 0 Activities\) | 179](#)
- [Add a Device to Routing Director | 188](#)
- [Move a Device to Production \(Day 1 and Day 2 Activities\) | 197](#)

Add Network Resource Pools and Profiles (Day -2 Activities)

Before a network architect plans to add devices, they must add the following to Routing Director.

1. Network resource pools to automatically assign IPv4 addresses, loopback addresses, and BGP cluster IDs. See ["Add Network Resource Pools" on page 676](#).
2. Device profiles to define device-level configurations such as loopback addresses, BGP groups, PCEP configuration, and so on.

A device profile is created based on the role of a device in the network. For example, a device might be a provider edge (PE) device or a metro router. For the PE device, you can configure BGP, IS-IS, and add tunnels in your device profile, whereas for the metro router, you can configure only BGP and IS-IS protocols.

See ["Add a Device Profile" on page 208](#).

3. Interface profiles to define the routing protocol configurations (OSPF, IS-IS, RSVP, and LDP) on the device. See ["Add an Interface Profile" on page 222](#).

An interface profile is created based on the role of the interface, for example core-facing or customer-facing.

4. Port profiles to define the aggregate Ethernet interfaces (MTU and Tagging) on the device. See ["Add a Port Profile" on page 226](#).

What's Next

A network planner uses device profiles, interface profiles, and port profiles in a network implementation plan to define the device configurations. The network planner can use the profiles in the plan as per their specific needs. That is, apply the profiles to all the devices and ports or to specific devices and ports.

To add the network implementation plan. See ["Add a Network Implementation Plan" on page 241](#).

RELATED DOCUMENTATION

[Device Onboarding Overview | 171](#)

[Device and Interface Profiles Overview | 202](#)

[Network Implementation Plan Overview | 230](#)

Prepare for Device Onboarding (Day -1 Activities)

Before a device is onboarded to Routing Director, a network planner must create a network implementation plan for onboarding the device. In the network implementation plan, you (network planner):

- Add devices and device interfaces.
- Assign device profiles and interface profiles to the devices and interfaces in the plan. The profiles define the general and protocol configurations for the devices and interfaces. You can also configure specific configurations (for example, IP address) that override the configurations in the profiles.
- Define the connections between devices added to the plan.
- Specify the SFPs (also called as pluggables or optics) to be installed in the device ports and cables to be used for connecting the ports. The information related to the SFPs and cables is displayed to the field technician for guiding the field technician to install the device at a site.

Apart from using the network implementation plan for applying configurations on a device during onboarding, you can use the network implementation plan to make any changes to the device or interface configurations even after the device is onboarded. For example, after a device is onboarded, at a later time, to upgrade the software on the device, you change only the software version of the device

in the network implementation plan and publish the plan. Routing Director automatically upgrades the software of the devices added to the plan.

If you assign a device to a network implementation plan, Routing Director starts monitoring the device, which enables it to track the device's performance and health right after the device is onboarded.

For information about adding the network implementation plan, see ["Add a Network Implementation Plan" on page 241](#).

What's Next

After you create the network implementation plan for a device, the device can be installed and onboarded to Routing Director. See ["Install and Onboard the Device \(Day 0 Activities\)" on page 179](#).

RELATED DOCUMENTATION

[Device Onboarding Overview | 171](#)

[Device and Interface Profiles Overview | 202](#)

[Device Onboarding Workflow | 173](#)

Install and Onboard the Device (Day 0 Activities)

A device can be installed and onboarded in one of the following two ways:

- A field technician, a user assigned the Installer role, installs and onboards devices to Routing Director as part of the Day 0 activities.
- A network administrator can onboard a device by using the adopt method after the field technician installs the device. See ["Add a Device to Routing Director" on page 188](#) for details.



NOTE: You must install and onboard the devices associated with a network implementation plan one at a time.

Prerequisites

The field technician UI in Routing Director enables you to access the network implementation plan of a device and get guidance to perform the Day 0 tasks.

Before you start onboarding a device to Routing Director, ensure that you have:

- A laptop that you can use to access the field technician UI.

- Internet connectivity on the laptop.
- The credentials to log in to Routing Director.
- Names of the organization and the site at which the device must be installed and onboarded.
- Unboxed the device and mounted it on the rack.

For details about installing a Juniper device, see the corresponding device Hardware Guide or [Quick Start Guide](#) at the [Documentation](#) site. To access the Hardware Guide or the Quick Start Guide of a device, on the homepage of the Documentation site, under **DOCUMENTATION BY CATEGORY**, click **View More > Device-Model** in the **Routing** section.

For details about installing a third-party vendor device, see documentation from the respective vendor.

You must also ensure that:

- If a firewall exists between Routing Director and the device, the firewall is configured to allow outbound access on TCP ports 443, 636, 2200, 6800, 4189, and 32,767.
 - Destination URL for LDAP: *ldap-server: 636*
 - Destination URL for NETCONF: *web-ui-ip-address: 2200*.
 - Destination URL for gNMI: *web-ui-ip-address: 32,767*.
 - Destination URL for Paragon Active Assurance Test Agent: *test-agent-virtual-IP-address: 6800*.
 - Destination URL to access Routing Director from a laptop or a desktop: *web-ui-ip-address: 443*.
 - PCE server *VIP address*.

Where, *ldap-server* is the IP address of the LDAP server, *web-ui-ip-address* is the IP address of the Routing Director Web UI, *test-agent-virtual-ip-address* is the Paragon Active Assurance virtual IP address, and *VIP-address* is the VIP address of the PCE server.

- Static routes are configured to reach the Internet.

```
user@device#set routing-options static route 0.0.0.0/0 next-hop Gateway-IP-address
```

- A DNS server is configured on the device to resolve domain names or the device is able to access an external DNS server (for example, 8.8.8.8).

To onboard a device to Routing Director:

1. Power on the device.
2. Log in to Routing Director by using the laptop.

The Select an organization page appears.

3. Choose the organization with which the device is to be associated.

By default, the Onboard a Device page appears.

4. Start the onboarding process in one of the following ways:

- On the Onboard a Device page, enter the serial number of the device.

If the serial number entered is correct, you can view the device name, device model, device serial number, and site (where the device is to be installed) for confirmation. Click **Yes** to confirm.

If the serial number entered is incorrect or if you do not know the serial number for a device, see ["Onboard a Device without a Serial Number" on page 185](#).

- Click Device List on the left navigation menu to view the list of devices to be onboarded on the Device List page. On the Device List page, the devices are listed with the serial number. Click the device you want to onboard.

5. Do one of the following:

- If a network implementation plan is present for the device, the Add a Device page appears.

The page provides information about the management interface of the device and any instructions (for example, **re0:mgmt-0.0: Insert QSFP28**) for inserting pluggables and connecting cables to the management interface.



NOTE: You can view the type of pluggable and the cable to be used for management port only if the instructions are added to the network implementation plan.

- If a network implementation plan does not exist, see ["Onboard a Device without a Network Implementation Plan" on page 186](#).

6. Insert the pluggable and connect the cable to the management interface as instructed in the network implementation plan.

7. Connect the management cable from the device to the network port.

8. Click **Adopt Device Manually** to view and copy the outbound SSH configuration to be committed on the device.

If the device is already connected to Routing Director by copying the outbound SSH command, go to step 19.

9. From the pop-up that appears, click **Copy**.

The outbound SSH configuration is copied to your clipboard.

10. Log in to the device by using SSH and enter the configuration mode.

11. Paste the contents of the clipboard and commit the configuration.

The device connects with Routing Director. Routing Director then starts pushing the device configuration and performing tests configured in the network implementation plan on the device. The Management Connectivity field turns green on the UI when Routing Director starts the tasks listed in the network implementation plan.

As the first test, Routing Director checks the model of the device. It then authenticates the device and verifies the validity of the software installed and displays the results.



NOTE: If the model of the device does not match the model of the device in the network implementation plan, the onboarding fails. Contact the network administrator for further action.

12. Click **Next** to view the Install Ports page.

On the Install Ports page, you can view the list of ports and the type of pluggables to be inserted in the ports.



NOTE: You can view the pluggables only for those ports for which instructions are added to the network implementation plan.

13. Insert pluggables as indicated on the Install Ports page.

After you insert the pluggables, Routing Director checks if the pluggable inserted matches the pluggable indicated in the network implementation plan for that port. If the pluggable inserted is correct, Routing Director tests the interface health of the pluggable.

14. Click **Next** to view the Install cables page.

On the Install Cables page, view the cabling instructions (if any provided) for the ports in the plan.



NOTE: You can view the instructions for connecting cables only for those ports for which instructions are added to the network implementation plan.

15. Connect Cables as indicated on the Install Cables page.

After you connect a cable, Routing Director initiates ping tests to the destinations defined in the plan.

16. Click **Next** to view the Testing page.

The Testing page shows the progress of all the tests being executed and the result of the tests. You can view the progress and the result of the tests listed in [Table 49 on page 183](#).

17. After onboarding completes:

- Onboarding successful is displayed if onboarding completes successfully.
- Onboarding failed is displayed if onboarding fails.

- Onboarding completed with errors is displayed if onboarding completes with errors.

Resolve any issues by using the failure reason provided or contact the Network Admin, if needed.

The onboarding status remains as failed even after you resolve the issues. The alerts, alarms, or warnings related to the issues continue to be displayed on the UI.

18. Click **Done.**

The device onboarding is complete.

19. Click **Device List to view the list of devices included in the network implementation plan.**

The Onboard section lists all the devices in the plan that are to be onboarded and the Completed section lists all the devices that have completed onboarding.

20. Repeat step 1 through step 19 to onboard all the devices in the plan.

After device onboarding is complete, the next step is to move the device to production. See ["Move a Device to Production \(Day 1 and Day 2 Activities\)" on page 197](#).

Table 49: Device Onboarding Tests

Test	Result Description
Management connectivity	You can view whether the device has established an outbound SSH connection to Routing Director. If the connection is not established, you must adopt the device manually. See "Add a Device to Routing Director" on page 188 to adopt the device.
Hardware health	You can view the health of fans, line cards, power supply modules (PSMs), CPU, and memory.
Interfaces	You can view the health of the interfaces.
Optics	You can view whether you have installed the correct pluggable in the port. The test result is displayed as follows: • Green check mark—Correct pluggable is inserted and the optics are healthy. • Yellow check mark—Either health data is not received from the pluggable or pluggable health check is disabled in the plan. • Red check mark—A wrong pluggable is inserted in the port or pluggable is missing.

Table 49: Device Onboarding Tests (*Continued*)

Test	Result Description
Chassis Alarm	You can view the severity of alarms present on the device. The following color code is used to display the chassis alarms present on the device: • Green—No alarms are present. • Red—At least one critical alarm is present. • Yellow—At least one minor alarm is present.
Trust Result	You can view whether the hardware is authentic and whether a valid software is installed on the device.
Neighbor Connectivity	You can view the result of the ping tests from the device port to neighboring devices. NOTE: You can view the results of the ping tests only for those ports for which the Active Assurance tests are enabled in the network implementation plan.
Remote Connectivity	You can view the results of the following connectivity tests from the device ports to the remote endpoints (edge devices, Internet endpoints, and cloud provider endpoints configured in the network implementation plan): • HTTP tests • DNS tests • Ping tests to the edge devices • Ping tests to the cloud provider endpoints NOTE: You can view the results of the connectivity tests only for those ports for which the Active Assurance tests are enabled in the network implementation plan.

Onboard a Device without a Serial Number

A network implementation plan can include one or more devices without a serial number. This is because, it is possible that while planning for device onboarding, the device is not yet procured and therefore the serial number is unavailable.

If the serial number entered on the Onboard a Device is incorrect, the UI provides a pop-up with Adopt and Select options. You (field technician) can:

- Click **Adopt** to commit the outbound SSH configuration for the device. To adopt the device:
 1. Click **Adopt** to view the outbound SSH commands that you can commit on the device to connect with Routing Director.

The Select Site pop-up window appears.

2. Select the site (location) to install the device.

3. Click **Adopt**.

A pop-up window displays the SSH commands that you must copy and commit on the device.

4. Click **Copy**.

The commands are copied to clipboard and the pop-up window closes.

5. Access the device in the configuration mode and commit the commands.

The device connects with Routing Director. The device is listed on the Inventory page (*device-type* tab) of the Routing Director UI with status as Connected.

The device can now be managed by using the element management functions in Routing Director. See ["Device Management Workflow" on page 355](#).

- Click **Select** to associate a serial number to the device.

A message appears for you to confirm whether you want to assign a serial number to the device. Do one of the following:

- Use the Onboard a Device page to proceed:
 1. On the Onboard a Device page, enter the serial number of the device in the **Serial Number** field.
 2. Click **OK**.

The Device List page appears. The Device List page displays the list of devices without a serial number under the Onboard section.

3. Click a device without a serial number in the Device List to associate the device with a serial number that you entered. The device that you select must match the model of the device you want to onboard.

The Add a Device page appears.

4. Go to step 5 of the onboarding procedure to proceed from the Add a Device page.
- Use the Device List page to proceed:

1. Click the Device List menu to open the Device List page.

The Device List page lists all devices to be onboarded along with their serial number. The page also displays a few devices without a serial number.

2. Click on a device that does not have a serial number and that matches with the model of the device you want to onboard.

You get a message asking to confirm whether you want to associate a serial number.

3. Click **OK**.

The Onboard a Device page appears.

4. Enter the device's serial number in the **Serial Number** field.

The Add a Device page appears.

5. Go to step 5 of the onboarding procedure to proceed from the Add a Device page.

Onboard a Device without a Network Implementation Plan

When a network implementation plan does not exist, you can:

- Adopt the device by committing the outbound SSH configuration.
1. Click **Adopt** to view the outbound SSH commands that you must commit on the device to connect with Routing Director.

The Select Site pop-up window appears.

2. Select the site (location) to install the device.

3. Click **Adopt**.

A pop-up window displays the SSH commands that you must copy and commit on the device.

4. Click **Copy**.

The commands are copied to clipboard and the pop-up window closes.

5. Access the device in the configuration mode and commit the commands.

The device connects with Routing Director. The device is listed on the Inventory page (Router tab) of the Routing Director UI with status as Connected.

The device can now be managed by using the element management functions in Routing Director. See ["Device Management Workflow" on page 355](#).

- Associate the serial number of the device with a device in a network implementation plan that does not have a serial number and use that plan for onboarding the device.

1. Click the **Device List** menu to open the Device List page.

The Device List page lists all the devices to be onboarded along with the serial number.

2. Click a device that does not have a serial number. The device you select must match the model of the device you want to onboard.

You get a message asking to confirm whether you want to associate a serial number.

3. Click **OK**.

The Onboard a Device page appears.

4. Enter the device's serial number in the **Serial Number** field.

5. Go to step 5 of the onboarding procedure to proceed from the Add a Device page.

What's Next

If the device is onboarded device by using the network implementation plan, a Super User or Network Admin can move the device to production and manage the device by using the network implementation plan. See ["Move a Device to Production \(Day 1 and Day 2 Activities\)" on page 197](#).

If the device does not use the network implementation plan to onboard a Super User or Network Admin can manage the device by using the element management functions in Routing Director. See ["Device Management Workflow" on page 355](#).

RELATED DOCUMENTATION

[Manage Device Licenses | 361](#)

[Deploy a Configuration Template to a Device | 390](#)

Add a Device to Routing Director

IN THIS SECTION

- [Adopt a Juniper Device | 189](#)
- [Adopt a Juniper Device Using ZTP | 190](#)
- [Add a Non-Juniper Device | 194](#)

You must be a superuser or network administrator to adopt or onboard a device (router, switch, or firewall).

A superuser or network administrator can adopt a device [both new devices (greenfield) and devices that are already a part of the network (brownfield device)], and manage the device by using Routing Director. When you adopt a device that is not associated with a network implementation plan, you (superuser or network administrator) must manually update configurations by using configuration templates, apply licenses, and upgrade software. However, if you use a network implementation plan to onboard devices, you can make the changes to the network implementation plan and publish the network implementation plan for the changes to take effect on the devices included in the plan. You also obtain the granular metrics about the device's health and performance by using the network implementation plan to onboard a device.

The status of a device that is already installed and connected to the network, but is not managed by the Routing Director appears as Disconnected on the Inventory page (**Inventory > Devices > Network Inventory**). When you adopt a device, the device connects with Routing Director and the status of the device changes to Connected, indicating that the device is managed by Routing Director.

You can add a device to Routing Director by using any of the following methods:

- Add a Juniper device; see ["Adopt a Juniper Device" on page 189](#).
- Adopt a Juniper device by using ZTP; see ["Adopt a Juniper Device Using ZTP" on page 190](#).
- Add a non-Juniper device; see ["Add a Non-Juniper Device" on page 194](#).

Before you adopt a device, ensure that:

- The device can reach the gateway.



NOTE: If a firewall exists between Routing Director and the device, configure the firewall to allow outbound access on TCP ports 443, 2200, 6800, 4189, and 32,767 from the management port of the device.

- The device can connect to Routing Director.

Adopt a Juniper Device

To adopt a device:

1. Navigate to **Inventory > Devices > Network Inventory**.
The Inventory page appears.
2. On the *device-type* tab, click **Add Device**.
The Add Devices page appears.
3. To add a Juniper Router, click **Adopt Router**.
4. (Optional) Click **Select Site** to select the site where the device is installed or to be installed.
5. In the **Select IP Version** drop-down list, select:
 - **IPv4** (default) to use IPv4 address for connecting with Routing Director.
 - **IPv6** to use IPv6 address for connecting with Routing Director.



NOTE: If Routing Director cluster uses IPv6 address, the outbound SSH command always uses IPv6 address irrespective of the IP version you select.

6. (Optional) If you are using RADIUS for device authentication, enter the IP address of the RADIUS server in the RADIUS Source Address field.



NOTE: Ensure that the version of the device IP address that you enter here matches with the version of the RADIUS server IP address that you have configured in Routing Director.

7. Click **Copy Cli Commands** to copy the CLI commands under **Apply the following CLI commands to adopt a Juniper Device if meets the requirements** and click **OK**.
The outbound SSH command listed depends on the IP version that you select in step 5.
8. Access the device by using Telnet or SSH and log in to the device in configuration mode.
9. Paste the contents of the clipboard and commit the configuration on the device.
The device connects to and is managed by Routing Director.

10. After you adopt a device, you can verify the connectivity status by running the following command on the device:

```
user@host> show system connections |match 2200
```

An output similar to the following indicates that the device is connected to Routing Director:

```
tcp 0 0 ip-address:38284 ip-address:2200 ESTABLISHED 6692/sshd: jcloud-s
```

What's Next

Connect to the Routing Director GUI and view the status of device onboarding. After the device status changes to Connected on the Inventory page (**Inventory > Devices > Network Inventory**), you can start managing the device. See ["Device Management Workflow" on page 355](#).


Adopt a Juniper Device Using ZTP

Prerequisites:

- A network implementation plan should be configured for the device.
- The device should be returned to zero or in its factory-default settings.
- A TFTP server reachable from the device
- A DHCP server reachable from the device, with the ability to respond to the device with the TFTP server and configuration file (script) name

[Figure on page 191](#) shows the workflow for adopting device using ZTP.

Figure 14: Workflow for Adopting a Device Using ZTP



The steps for the workflow are as follows:

1. Create an onboarding script (Python or SLAX) with the required SSH configuration by:
 - Connecting to Routing Director GUI and copying the onboarding configuration statements. To copy the onboarding configuration from the Routing Director GUI:
 - a. Navigate to **Inventory > Network Inventory**.
 - b. On the *device-type* tab, click **Add Device**.
 - c. On the Adopt Router tab of the Add Devices page,
 - i. (Optional) Click **Select site** to select the site where the device is installed.
 - d. Click the **Copy Cli Commands** link to copy the outbound SSH configuration displayed in the **Apply the following CLI commands to adopt a Juniper Device if meets the requirements** text field to clipboard and click **OK**.
 - e. Paste the CLI commands in the onboarding script.
 - Adding the SSH configuration statements to the onboarding script.
- See ["Sample Onboarding Script for Committing SSH Configuration on a Device "](#) on page 192 for a sample of the onboarding script.
2. Upload the onboarding script to the TFTP server.

3. Configure the DHCP server with the onboarding script filename and path to the onboarding script in the TFTP server.
4. Install the device, connect it to the network, and power on the device.
For information about installing the device, see the respective installation guide on the [Product Documentation](#) website.
5. After the device is powered on, the factory-default settings in the device trigger a built-in script (**ztp.py**). The script obtains the following from the DHCP server:
 - IP address for the management interface
 - IP address of the default gateway
 - IP address of the DNS server and TFTP server
 - Path to the onboarding script on the TFTP server
6. The device configures its management IP address, static default route, and the DNS server address, based on the values from the DHCP server.
7. The device downloads the onboarding script, based on the values from the DHCP server, and executes it, resulting in the onboarding configuration statements to be committed.
8. The device opens an outbound SSH session with Routing Director based on the committed onboarding configuration.
9. Routing Director configures management and telemetry parameters including gNMI by using NETCONF. Routing Director also configures the interfaces and protocols based on the configurations defined in the network implementation plan associated with the device, by using NETCONF.

Sample Onboarding Script for Committing SSH Configuration on a Device

The following is a sample of the onboarding script that is downloaded from the TFTP server to the device:

```
#!/usr/bin/python
from jnpr.junos import Device
from jnpr.junos.utils.config import Config
from jnpr.junos.exception import *
import sys

def main():
    config = "set system services ssh protocol-version v2\n\
set system authentication-order password\n\
set system login user jcloud class super-user\n\
set system login user jcloud authentication encrypted-password"
```

```

$6$0i4IvHbbFYT.XgXP7$43TeEU7V0Uw3CB1N/
HFKQT.Xl2wsm6GEBaS9pfE9d3VrINIKBqlYlJfE2cTcHsCSSVboNnVtqJEaLNUBAfbu.\n\
set system login user jcloud authentication ssh-rsa \"ssh-rsa
AAAAB3NzaC1yc2EAAAADAQABAAQgQCuVTPiYaDwBuB8aTVrzxDQ050BS5GtoGnMBkWBiyi5EEc0n8eJGmbINE8auRGGOtY/
CEbIHKSp78ptdzME0uQhc7UZm4Ue18C3FRb3qEYjr1AMJMU+hf4L4MYWYXqk+Y9RvnWBzsT02iEqGU0Jk0y7Mev5z/
YI9r8u8MZlWKdQzegBRiKL4HYy0AeAbenNw6ddxRzAP1bPESpmsT+0kChu3jYg8dzKbI
+xjDBhQsKCFf05cXyALjBmI3beaxmXRV02UGCEB1+5Xw6a30CiP7jplr92rFBjbqgh/bYoJRYz1Rc3AirDjR0QuDdpHRn
+DuUjPlyV17QR9Qvwn40AmWM9YKWS/
LZ375L8nac0Hmlv4f0KETU4LScTFQXR6xiJ6RizEp0338+xmiVq6m0cv5VuXfNAPd18F3LW0xLGF1mieB4cEEyJ7MK9U
+TgS7MlCAP+XAeXYM2Vx1b+UCyYoEyDizaRXZvmP5BPpxpb5L2iuXencZMbbpEbnNX/sk3teDc=
jcloud@5c96fb73-4e3a-4d8b-8257-7361ef0b95e7\"\\\"\\n\
set system services outbound-ssh client jcloud secret
f72b785d71ea90017f911a5d6c908c95f12a265e19e886f07a364ce12aa99c6c1ca072a1ccc7d39b3f8a7c94e7da761d1
396714c0b32ef32b6e7d3c9ab62cf49d8d\\n\
set system services outbound-ssh client jcloud services netconf keep-alive retry 12 timeout 5\\n\
set system services outbound-ssh client jcloud oc-term.cloud.juniper.net port 2200 timeout 60
retry 1000\\n\
set system services outbound-ssh client jcloud device-id
5c96fb73-4e3a-4d8b-8257-7391ef0b95e7.0ad21cc9-1fd6-4467-96fd-1f0750ad2678\\n\
set system root-authentication encrypted-password \"$6$0eRp2LWC$/
ZLm9CMiR.SeEunv.5sDksFHIkzafuHLf5f7sp1ZANYT0iiz6rk2A1d/4Bq1gmxBhEb1XFtskrocLD7VHvPU10\\\"\"
    dev = Device()
    dev.open()
    try:
        with Config(dev, mode=\"exclusive\") as cu:
            print (\"Loading and committing configuration changes\")
            cu.load(config, format=\"set\", merge=True)
            cu.commit()
    except Exception as err:
        print (err)
    dev.close()

if __name__ == \"__main__\":
    main()

```

What's Next

Connect to the Routing Director GUI and view the status of device onboarding on the Inventory page (**Inventory > Devices > Network Inventory**). After the device status changes to Connected, you can start managing the device. See ["Device Management Workflow" on page 355](#).

Add a Non-Juniper Device



NOTE: In this release, among non-Juniper devices, only Cisco Systems devices are supported. For a list of supported Cisco Systems devices, see [Supported Junos OS Releases, Devices, and Browsers](#).

To add a non-Juniper device to Routing Director:

1. Navigate to **Inventory > Devices > Network Inventory**.

The Inventory page appears.

2. On the Router tab, click **Add Device**.

The Add Devices page appears.

3. On the Add a Device tab, enter values referring to table [Table 50 on page 194](#).

4. Click **+ Add Device** to add more devices.

5. Repeat steps 3 and 4 to add more non-Juniper devices.

6. Click **OK**.

The device connects to Routing Director. You can now manage the device by using Routing Director.

7. (Optional) View the status of device onboarding on the Inventory page (**Inventory > Devices > Network Inventory**). After the device status changes to Connected, you can start managing the device. See "[Device Management Workflow](#)" on page 355.

Table 50: Add a Non-Juniper Device

Field	Description
Device Name	Enter the device hostname.
Address/GRPC Port	Enter the IPv4 address and gRPC port configured on the device.
MAC	Enter the MAC address of the device.
Site	Click the drop-down list to select a site where the device is deployed. Selecting a site to onboard a device is optional.
Vendor	Click the drop-down list to select the vendor of the device.
Model	Enter the device model. For example, Cisco NCS-57C3,

Table 50: Add a Non-Juniper Device (*Continued*)

Field	Description
Operating System Name	Enter the name of the operating system (OS) deployed on the device. For example, IOS XR.
Version	Enter the version of the operating system. For example 7.10.1.
Timeout (in minutes)	Enter the time (in minutes) of inactivity after which the device disconnects with Routing Director.
Retry Delay (in minutes)	If the device is unable to connect with Routing Director, enter the time (in minutes) after which the device retries to connect with Routing Director.
Authorization	
Insecure	This option is disabled by default and the connection between the device and Routing Director is secure (encrypted and/or authenticated). Enable this option when TLS is disabled on the device so that the connection between the device and Routing Director is established without any encryption. The insecure connection is vulnerable to man-in-the-middle attacks and other security threats. We recommend that you use an insecure connection between the device and Routing Director only during local development and in testing environments where security is not a concern.

Table 50: Add a Non-Juniper Device (*Continued*)

Field	Description
Skip Verify	This option is disabled by default and Routing Director verifies the device when the device establishes a connection with Routing Director. Enable this option when TLS is enabled on the device and Routing Director should skip verifying the device's identity when the device establishes a connection. This option is useful when you have a self-signed certificate that cannot be verified against a certificate authority, so that device verification is skipped. When you skip verifying the device's identity, the connection is vulnerable to man-in-the-middle attacks. Therefore, this option should be enabled only during development or testing the connection between the device and Routing Director.
Certificates—Use this option to enter the username and password for the device to authenticate with Routing Director.	
Certificate	Click Browse to upload the TLS certificate for the device.
Key Certificate	Click Browse to upload the TLS certificate key for the device.
Certificate Authority	Certificate Authority: Click Browse to upload the TLS root certificate from the certificate authority.
Credentials—Use this option to enter the username and password for the device to authenticate with Routing Director.	
Username	Enter a username of the device for Routing Director to connect with the device.
Password	Enter the password for the username.

SEE ALSO

[Device Onboarding Overview](#) | 171

[Device Management Workflow](#) | 355

Move a Device to Production (Day 1 and Day 2 Activities)

After a device is onboarded, a Network Operation Center (NOC) engineer (a user with the Network Admin or Super User role) can perform the following tasks before moving the device to production:

- View the list of devices that are ready for service.
- Apply additional configurations, if needed, for moving the device to production. See "[Deploy a Configuration Template to a Device](#)" on page 390.

You can apply additional configurations by:

- Editing a network implementation plan that you used to onboard the device. See "[Edit a Network Implementation Plan](#)" on page 260.
- Using configuration template. See "[Deploy a Configuration Template to a Device](#)" on page 390.
- On the *Device-Name* page (Inventory > Onboarding Dashboard > *Device-Name*), monitor and analyze the performance of the onboarded device.

To move a device to production:

1. Log in to Routing Director and access your organization.
The Troubleshoot Devices page appears.
2. On the left navigation menu, click **Inventory > Onboarding Dashboard**.
The Put Devices into Service page appears. The page displays a summary of the number of devices that are ready to be installed and ready for service along with a summary of the number of devices that have critical (Urgent Action Needed) and major (Action Needed) alerts and alarms.
3. Filter the devices that are ready for Service by selecting **Ready for Service** in the **Operational State** filter.
The devices with the Ready for Service status are listed.
4. Click the *Hostname* link of the device to view the result of the automated tests that are performed on the device.
The *Device-Name* page appears.
5. Analyze the results of the tests in the different accordions and view the alerts raised for the device.
If no critical or major issues are present, you can move the device to production.
6. Click **Put into Service** to move the device to production.

Routing Director changes the status of the device to In Service and moves the device to production.

7. You can configure services on the device and monitor the performance of the device from the *Device-Name* page. See ["Service Provisioning Workflow" on page 648](#).

RELATED DOCUMENTATION

[Device Management Workflow | 355](#)

[Service Orchestration Overview | 646](#)

Field Technician User Interface

IN THIS CHAPTER

- [Field Technician UI Overview | 199](#)
- [Working with Field Technician UI Pages | 200](#)

Field Technician UI Overview



NOTE: The Field Technician User Interface (UI) is an experimental feature.

Routing Director provides a field technician UI to view the list of devices that you (field technician) can onboard and start the onboarding process. You can access the field technician UI by using any handheld device such as a laptop. The UI provides an option to enter the serial number of the device to start the onboarding process. After you enter the serial number, the UI guides you to install pluggables and connect cables to the ports based on instructions entered in the network implementation plan.

If a network implementation plan is present for a device, the device automatically connects with Routing Director and provides instructions for inserting cables and pluggables. Routing Director also executes compliance and health checks on the device. The compliance checks must be enabled for the device in the plan used to onboard the device.

If an implementation plan does not exist for a device, the field technician UI provides:

- The outbound SSH commands that you can commit on the device to connect with Routing Director.
- An option to obtain a network implementation plan by assigning the device's serial number to a device in a plan that does not have a serial number and use that network implementation plan. For example, if a device D1 is not associated with any plan, then D1's serial number can be assigned to device D2 associated with a plan P1. D2 does not have a serial number assigned to it in the plan P1.

See ["Onboard a Device without a Network Implementation Plan" on page 186](#).

Routing Director checks for the interface health as you insert the pluggables and connect cables, and flags issues, if any. You can click the Expand icon on each row of the UI pages to view the details of the

checks. You can correct the errors and use the **Resume Onboarding** option to trigger Routing Director to resume testing from where the onboarding process paused. Besides testing the interface health, Routing Director also checks the health of the management connection and connection with neighbors, health of chassis components (fans, power supply modules (PSM), memory, line cards, and CPU). Routing Director also runs compliance scans to determine the authenticity, vulnerabilities, and trustworthiness of the device.

At the end of the onboarding process, the UI displays the results of all the tests and flags any issues that Routing Director has detected.

You must be assigned the Installer role to access the field technician UI. The field technician UI displays the following pages:

- Onboard a Device for onboarding devices by entering a device's serial number.
- Device List for viewing the list of onboarded devices and the list of devices that must be onboarded.

See "[Working with Field Technician UI Pages](#)" on [page 200](#) for more information.

RELATED DOCUMENTATION

| [Install and Onboard the Device \(Day 0 Activities\)](#) | 179

Working with Field Technician UI Pages

IN THIS SECTION

- [Onboard a Device Page](#) | 201
- [Device List Page](#) | 201

Routing Director provides a dedicated UI to guide a field technician to onboard devices. The UI displays the following two pages:

- Onboard a Device

Use this page to access instructions for installing pluggables and connecting cables to device ports and initiate the device onboarding process. You (field technician) can enter the serial number of the device to start the onboarding process.

- Device List

Use this page to view the list of devices that are already onboarded and the list of all the devices that you can onboard. Devices configured in the network implementation plans appear in the list of devices that you can onboard.

Onboard a Device Page

To access the Onboard a Device page:

1. Log in to Routing Director as an Installer.
The Select an Organization page appears.
2. Click the organization with which you want to associate devices.
The Onboard a device page appears.
3. Enter the serial number of the device that you want to onboard. Follow the instructions on the UI to onboard the device. See ["Install and Onboard the Device \(Day 0 Activities\)" on page 179](#).
4. (Optional) Refer to the Device List page to see if you must onboard more devices. See ["Device List Page" on page 201](#).
5. (Optional) Repeat steps 1 through 4 to onboard all the devices.

The onboarding process is complete after you complete onboarding all the devices.

Device List Page

To access the Device List page:

1. Log in to Routing Director as an Installer.
The Select an Organization page appears.
2. Click the organization with which you want to associate devices.
The Onboard a device page appears.
3. Click **Device List**.
The Device List page appears. Under the:
 - **Onboard** section, you can view the list of devices that you can onboard.
 - **Completed** section, you can view the list of devices that have completed onboarding.
4. Click a device in the Onboard section to start the onboarding process. See ["Install and Onboard the Device \(Day 0 Activities\)" on page 179](#).

SEE ALSO

[Device Onboarding Overview](#) | 171

[Device Onboarding Workflow](#) | 173

Onboarding Profiles

IN THIS CHAPTER

- [Device and Interface Profiles Overview | 202](#)
- [About the Device and Interface Profiles Page | 205](#)
- [Add Labels | 207](#)
- [Add a Device Profile | 208](#)
- [Add an Interface Profile | 222](#)
- [Add a Port Profile | 226](#)
- [Edit and Delete a Label or Profile | 227](#)

Device and Interface Profiles Overview

Onboarding profiles comprise device and interface profiles. A device onboarding profile defines the following parameters associated with a device:


- IP address
- Protocols—BGP, IS-IS, and OSPF
- Traffic engineering (TE)
- Segment routing (SR)
- Neighboring devices, DNS servers, Internet endpoints, and cloud provider endpoints for checking connectivity.
- Placement resources for reserving resources (bandwidth, MAC addresses, and routes) to provision service.
- Routing protocol analytics to enable BGP and IGP data collection.

A network architect creates a device profile based on the role of the device in the network. For example, a device might be a provider edge (PE) device or a metro router. For the PE router, you can configure BGP, IS-IS, and add tunnels in your device profile, whereas for a metro router, you can configure only the

BGP and IS-IS protocols. See ["Add a Device Profile" on page 208](#) for information about adding a device profile.

[Figure on page 203](#) shows the Create Device Profile page.

Figure 15: Add a Device Profile Page



An interface profile defines the configuration for OSPF, IS-IS, LDP, and RSVP protocols for an interface. You create an interface profile based on the role of an interface. For example, core-facing or customer-facing. See ["Add an Interface Profile" on page 222](#) for information about adding an interface profile.

[Figure 16 on page 203](#) shows the Create Interface Profile page.

Figure 16: Add an Interface Profile Page



You assign the profiles to a network implementation plan. In the network implementation plan, you can specify the profile that must be applied to all the devices or ports or specific devices or ports. Routing Director commits the configurations defined in the profiles and in the plan during device onboarding so that the device is ready for service soon after you onboard the device. You can also use the profiles to modify the configurations after a device is onboarded and in service. See ["Network Implementation Plan Overview" on page 230](#) for details about the network implementation plan.

In addition to device profiles and interface profiles, you can create port profiles for configuring aggregated Ethernet interfaces on a device. In a port profile, you can specify the MTU and specify whether the port would be VLAN tagged (trunk port) or untagged (access port). When you assign a port profile to a device's aggregated Ethernet interface port, the MTU and port type is configured on the aggregated Ethernet interfaces of the device to which the port profile is assigned when the device is onboarded to Routing Director. Later, to modify the MTU and the port type of the assigned aggregated Ethernet interface, you can modify the port profile. [Figure 17 on page 204](#) shows the Create Port Profile page. For information about adding a port profile, see ["Add a Port Profile" on page 226](#)

Figure 17: Port Profile

Create Port Profile

✕

Profile Name *

Plan Name

MTU

Tagging

Please select option(s) ▾

Cancel

Save

RELATED DOCUMENTATION

Add Network Resource Pools and Profiles (Day -2 Activities)	177
Device Onboarding Workflow	173

About the Device and Interface Profiles Page

IN THIS SECTION

- [Tasks You Can Perform](#) | 205

To access this page, click **Inventory > Device and Interface Profiles** on the left navigation menu.

You can use device and interface profiles to define the configurations that you want to commit on a device. The profiles include device-level parameters such as IP address assignment, autonomous system (AS) number, and protocol configurations. These configurations enable the device to be ready for service soon after the device is onboarded. After a device is onboarded, you can use the profiles to modify the configurations as well.

Tasks You Can Perform

- View details of device and interface profiles—The Profiles table displays the various parameters of the configured device and interface profiles. See [Table 51 on page 206](#).
- Add labels; see ["Add Labels" on page 207](#).
- Add a device profile; see ["Add a Device Profile" on page 208](#).
- Add a port profile; see ["Add a Port Profile" on page 226](#).
- Add an interface profile; see ["Add an Interface Profile" on page 222](#).
- Edit and delete device, interface, and port profiles; see ["Edit and Delete a Label or Profile" on page 227](#).
- Search by using keywords—Click the search icon (magnifying glass), enter the search term in the text box, and press Enter. The search results are displayed on the same page.
- Filter the data displayed in the table—Click the filter icon (funnel) and select whether you want to show or hide advanced filters. You can then add or remove filter criteria, save criteria as a filter, apply or clear filters, and so on. The filtered results are displayed on the same page.
- Show or hide columns in the table or reset page preferences, using the vertical ellipsis menu.
- Sort, resize, or re-arrange columns in a table (grid).

- Set the number of rows to be displayed in the table—Set the number of rows (15, 30, 50, or 100) to be displayed in the Network Implementation Plan table in the **Display** drop-down list provided at the bottom right-corner of the table.

If there are multiple pages of rows, use the arrow heads to traverse through the pages or enter a page number to view the entries listed on that page.

Table 51: Fields on the Device and Interface Profiles Page

Field	Description
Type	Label or type of profile—Device, Interface, or port. A label is a keyword for referencing a group of devices.
Name	Name of the label or profile.
Plan Name	Name of the network implementation plan to which the profile is assigned. A network implementation plan defines the assignment of device and interface profiles to one or more devices.
Protocols	Protocols configured in the profile. By default, this column lists two protocols configured in the profile. If you configure more than two protocols, <i>+integer</i> (for example: +2) appears to the right of the protocol. The integer indicates the total number of protocols configured in the profile. Click the integer to view all the protocols defined in the profile.
Labels	Labels assigned to the profile.

RELATED DOCUMENTATION

[Add Network Resource Pools and Profiles \(Day -2 Activities\) | 177](#)

[Add Network Resource Pools | 676](#)

Add Labels

A label references a group of devices. You use labels to identify devices of the same type or role and reference them in a device profile.

For example, you can tag all provider edge devices with the label PE. Then, within a device profile, you can define that BGP sessions or MPLS LSPs should be established with any other device with the same label. When a provider edge device is onboarded using this profile, it gets tagged with label PE and automatically configured to peer with all the other devices also tagged with the label PE. At the same time, all these other devices also get configured to peer with this new device.

You associate a label with a network implementation plan.

To add a label:

1. Navigate to **Inventory > Devices > Device and interface Profiles**.

The Device and Interface Profiles page appears.

2. Click **Add > Labels**.

The Create Labels page appears.

3. Enter values by referring to [Table 52 on page 207](#).

Fields marked with * are mandatory.

4. Click **OK**.

The label is created and listed on the Device and Interface Profiles page.

Table 52: Fields on the Add Labels Page

Field	Description
Plan Name	Enter a name for the network implementation plan with which the label is associated. The name can contain alphanumeric characters and some special characters [hyphen (-) and period (.)] and cannot exceed 64 characters.
Label*	Enter one or more labels to be associated with the network implementation plan. A label is used to reference a group of devices (instead of referencing individual devices) in a profile. For example, in a device profile, you can enter the device label to specify devices to be included in a BGP peer group instead of entering IP addresses of individual devices. All devices that are associated with the label become part of the same BGP peer group. The name of a label can contain alphanumeric characters and some special characters [hyphen (-), underscore (_), and period (.)] and cannot exceed 64 characters.

RELATED DOCUMENTATION

Add a Network Implementation Plan 241
Device Onboarding Overview 171

Add a Device Profile

A device profile defines global configuration elements that are added to the device during onboarding. The configuration elements include hostname, IP address of the loopback, router ID, AS number, routing protocols, segment routing, bandwidth, MAC addresses, end points for testing connectivity, traffic engineering, and routing protocol analytics.

We recommend that you create some device profiles with configurations that can be applied to all the devices in a network implementation plan and some profiles with device-specific configurations.

Before you create device profiles, ensure that you have the required network resource pools (for example, IP addresses and BGP cluster IDs) configured in Routing Director. If you configure Routing Director to assign values for network resources (loopback addresses, IPv4 addresses, BGP cluster IDs, and so on), Routing Director uses the network resource pools to assign the values. See ["Add Network Resource Pools and Profiles \(Day -2 Activities\)" on page 177](#). for details.

To add a device profile to Routing Director:

1. Navigate to **Inventory > Devices > Device and Interface Profiles**.
The Device and Interface Profiles page appears.
2. Click **Add > Device Profile** to create a device profile.
The Create Device Profile page appears.
3. Enter values by referring to [Table 53 on page 208](#).
Fields marked with * are mandatory.
4. Click **Save** to save the profile.
You can view the profile listed on the Device and Interface Profiles page.

Table 53: Fields on the Create Device Profile Page

Field	Description
General	

Table 53: Fields on the Create Device Profile Page (*Continued*)

Field	Description
Upload JSON File	Click Browse to upload a pre-created device profile in the JSON file format. The values in the pre-created device profile are automatically populated in the Create Device Profile page. Click the Download this form into JSON file link to download and to save the profile in its current state (for example, when you want to save the current configured values for later reference or for maintaining a record).
Profile Name*	Enter a name for the device profile. The profile name can contain alphanumeric characters and some special characters [hyphen (-), underscore (_), period (.), and colon (:)] and cannot exceed 64 characters.
Plan Name	Enter a name for the network implementation plan in which you want to use this profile. You can use the device profile only in the network implementation plan that you enter here. A network implementation plan with the name you enter here is auto-generated and listed on the Network Implementation Plan page (Inventory > Device Onboarding > Network Implementation Plan). The plan name can contain alphanumeric characters and some special characters [hyphen (-) and period (.)] and cannot exceed 64 characters.
Configuration Template	Select one or more configuration templates that you want to include in the profile. Use configuration templates to configure parameters that you cannot define in the profile or plan. All the configuration templates present in the organization are listed here. NOTE: The configurations in the configuration template can be applied only during onboarding. Any changes to the configurations made after the device is onboarded are not applied to the device.

Table 53: Fields on the Create Device Profile Page (*Continued*)

Field	Description
Device Labels	Select one or more device labels from the drop-down list. The labels that you select here are associated with the devices to which you assign this profile. You can use the labels to refer to the device in various contexts. For example, if you assign the label PE for provider edge devices, you can use the label to filter all PE devices present in your network. You can also click the + Add new label link to add a new label to the profile, in the Add New Label page. The name of the label can contain alphanumeric characters and some special characters [hyphen (-) and period (.)] and cannot exceed 64 characters. See "Add Labels" on page 207.
Software Image	Select the software image to be installed on the device. During device onboarding, Routing Director checks whether the software version installed on the device matches the version you enter here. If the software version does not match, the software version that you specify here is installed on the device. You can view the images that are uploaded to Routing Director here.
Autonomous System	Enter the ID or number of the AS to which you want to assign the device. Range: 1 through 4,199,999,999
Trust	Click to enable (default) or disable Routing Director to run compliance scans on the device for assessing the integrity and potential vulnerabilities on the device and to calculate compliance score for the device. The compliance score of a device indicates compliance of the device with the rules defined in the Center for Internet Security (CIS) benchmarks.
Router ID	Click to enable or disable (default) automatic router ID configuration on a device during device onboarding. If you enable automatic router ID configuration, the IPv4 loopback address of the device is used as the router ID.

Table 53: Fields on the Create Device Profile Page (*Continued*)

Field	Description
IPv4 Loopback Address	Click to enable or disable (default) automatic IPv4 loopback address configuration on the device. If you enable automatic IPv4 loopback address configuration, Routing Director assigns the IPv4 loopback address automatically to the device's interfaces from the IPv4 address resource pool. For automatic configuration of IPv4 loopback address, you must have IPv4 loopback address resource pools uploaded to Routing Director. Otherwise, the IPv4 loopback address is not assigned to the device's interfaces and device onboarding fails. See "Add Network Resource Pools" on page 676 for adding information about resources pools. If you disable this option, you can configure the loopback address when you add devices to a network implementation plan.
IPv6 Loopback Address	Click to enable or disable (default) automatic IPv6 loopback address configuration on the device. If you enable automatic IPv6 loopback address configuration, Routing Director assigns the IPv6 loopback address automatically to the device's interfaces from the IPv6 address resource pool. For automatic configuration of IPv6 loopback address, you must have IPv6 loopback address resource pools uploaded to Routing Director. Otherwise, the IPv6 loopback address is not assigned to the device's interfaces and device onboarding fails. See "Add Network Resource Pools" on page 676 for adding information about resources pools. If you disable this option, you can configure the loopback address when you add devices to a network implementation plan.
Per Flow Load Balancing	Click to enable (default) or disable load-balancing traffic over two or more paths when equal-cost multiple paths exist to a destination. For information about configuring per-flow load balancing, see Configuring per-flow load balancing.

Table 53: Fields on the Create Device Profile Page (*Continued*)

Field	Description
IPv6 tunneling	Click to enable or disable (default) MPLS LSPs to be used for IPv6 traffic. When you enable this option, LDP and RSVP routes in inet.3 routing table are converted IPv6 addresses (mapped to IPv4 addresses) and stored in inet6.3 routing table. This option is available only with infrastructure service design version 0.13.20 or later. See "Update the Network Implementation Plan After Routing Director Upgrade" on page 261 to upgrade your infrastructure service design version.
ISO Network Address	Click to enable or disable (default) IS-IS protocol configuration on the device. If you enable ISO Network Address, configure the area ID and system ID.
Area ID*	Enter the area ID to be assigned to the device for IS-IS protocol configuration. Range: 01 through 99
System ID	Click to enable (default) or disable auto-generation of a system ID for IS-IS protocol configuration. If you choose to auto-generate the system ID, the value assigned is usually the host part of the device's IPv4 or IPv6 loopback address in the binary-coded decimal (BCD) format. For automatic configuration of System ID, you must have IPv4 loopback address resource pools uploaded to Routing Director. Otherwise, the System ID is not assigned to the device and device onboarding fails. See "Add Network Resource Pools" on page 676 for information about adding resources pools. If you explicitly specify the system ID, we recommend that you use the IPv4 loopback address represented in the BCD format. For example, if the loopback address is 192.168.1.77, the system ID should be 1921.6800.1077.
Routing Protocols	
BGP	Click to enable or disable (default) BGP configuration on the device. If you enable BGP configuration, add an internal or external BGP peer group for the device. For information about the configurable fields to add a BGP group, See Table 54 on page 217. You can also edit and delete BGP peer groups of a device from here.

Table 53: Fields on the Create Device Profile Page (*Continued*)

Field	Description
PCEP	Click to enable or disable (default) path computation element protocol (PCEP) configuration on a device. If you enable PCEP, configure the IPv4 path computation element (PCE) address in your network.
PCE Address	IPv4 or IPv6 address of the PCE in your network.
Traffic Engineering	Click to enable or disable (default) traffic engineering (TE) configuration on your device. If you enable TE, add tunnels [label-switched paths (LSPs)] for TE. See Table 55 on page 220. You can also edit and delete tunnels from here. NOTE: If you configure tunnels, you must configure RSVP in an interface profile and apply the interface profile to a device to which you apply this device profile.
Segment Routing	Click to enable or disable (default) segment routing configuration on a device. If you enable segment routing, configure start label and index range for the OSPF and IS-IS protocols, and the node segment identifier (SID) (referred to as IPv4 index) for a device.
OSPF	
Start Label*	Enter a start label for the segment routing label block. This label is advertised using the OSPF protocol. Range: 16 through 1,048,575
Index Range*	Enter the range of label values that you want to use as the SID for a device. Range: 32 through 1,048,559
ISIS	

Table 53: Fields on the Create Device Profile Page (*Continued*)

Field	Description
Start Label*	Enter a start label for the segment routing label block. This label is advertised using the IS-IS protocol. Range: 16 through 1,048,575
Index Range*	Enter the range of label values that you want to use as SID for a device. Range: 32 through 1,048,559
IPv4 Index	Click to enable or disable (default) the automatic configuration of the IPv4 node SID for segment routing. For automatic configuration of IPv4 index, you must have the segment identifier resource pools uploaded to Routing Director. Otherwise, the IPv4 index is not assigned to the device and the device onboarding process fails. See "Add Network Resource Pools" on page 676 for information about adding resources pools.
IPv6 Index	Click to enable or disable (default) the automatic configuration of the IPv6 node SID for segment routing. For automatic configuration of IPv6 index, you must have the segment identifier resource pools uploaded to Routing Director. Otherwise, the IPv6 index is not assigned to the device and the device onboarding process fails. See "Add Network Resource Pools" on page 676 for information about adding resources pools.
ISIS	
IPv4	Click to enable (default) or disable IPv4 routing by using IS-IS protocol.
IPv6	Click to enable (default) or disable IPv6 routing by using IS-IS protocol.
LDP	
IPv4	Click to enable (default) or disable IPv4 routing by using LDP protocol.
IPv6	Click to enable (default) or disable IPv6 routing by using LDP protocol.

Table 53: Fields on the Create Device Profile Page (*Continued*)

Field	Description
Active Assurance	
Edge Devices	Click to enable or disable (default) the test agents installed on ACX routers and x86 platforms to run connectivity test to the edge devices in your network. If you enable running connectivity tests to the edge devices, configure the labels and IPv4 addresses of the edge devices.
Device Labels	Select the device labels for edge devices. Test agents run connectivity tests to all devices that share the device label.
Addresses	Enter the IPv4 addresses of edge devices to which test agents on the device run connectivity tests.
Internet Endpoints	Click to enable or disable (default) the test agents that are installed on devices to run connectivity tests to the Internet endpoints such as Web servers and DNS servers in your network. If you enable running connectivity tests to the Internet endpoints, you must configure the endpoints for the connectivity test.
Endpoints	Click + to add Internet Endpoints for connectivity checks. Configure the following: • Name—Enter the name of the Internet endpoint server. • URL—Enter the URL of the Internet endpoint server in host[:port]/[path] format. For example, www.example.com/v1. • Click Add common endpoints to select common endpoints from the list. Click the check mark to save the endpoints.
DNS Server	Enter the IPv4 address of the internal or external DNS server to which the test agent runs a ping connectivity test.

Table 53: Fields on the Create Device Profile Page (*Continued*)

Field	Description
Cloud Providers	Click to enable or disable (default) the test agents installed on devices from running connectivity tests to hosts in the Cloud Provider's network. If you enable running connectivity tests to the cloud provider endpoints, you must configure the cloud provider endpoints.
Select cloud providers	Configure the parameters to check connectivity from a device to the cloud provider network. To configure connectivity tests to cloud provider endpoints: 1. Select a cloud provider (Amazon Web Services [AWS], Microsoft Azure, or Google Cloud Platform) in the Cloud Providers list to which connectivity is to be tested. 2. (Optional) Click Edit to change the default delay and delay variance threshold values for the selected cloud provider. You can edit the values as per your preference and click the check mark to save the edited values. 3. Click Save. Routing Director runs connectivity checks to the configured cloud provider endpoints during device onboarding.
Placement Resources	
Bandwidth*	Enter the bandwidth (in Mbps) to be used for service provisioning on the device.
MAC Address	Enter the maximum number of MAC addresses to be assigned to the device for service provisioning.
Routes	Enter the maximum number of routes addresses to be assigned to the device for service provisioning.
Analytics	

Table 53: Fields on the Create Device Profile Page (*Continued*)

Field	Description
Traffic-Engineering	Click to enable or disable (default) the collection of historical telemetry data for the elements (tunnels, interfaces, delay, and so on) that are configured for traffic data collection.
Blackhole Detection	Click to enable (default) or disable Routing Director from displaying predictor events. Predictor events are kernel (OS), routing, and forwarding exceptions that can potentially cause traffic loss.
Routing Protocol Analytics	Click to enable or disable routing protocol analytics and data collection (in the device profile) that you can apply to devices in the network implementation plan. Routing protocol analytics are disabled by default. You can use the Route Explorer dashboard and its features only after you have enabled Routing Protocol Analytics.
BGP State	Click the check box to enable the data collection and monitoring of all BGP peers. You can view BGP state and analytics on the Route Explorer dashboard. See "About the Route Explorer Page" on page 566.
BGP Ribs	Click the check box to enable the data collection and monitoring of all BGP-related routing information bases (RIBs). You can view a record of all BGP RIBs on the Route Explorer dashboard. See "About the Route Explorer Page" on page 566.

Table 54: Fields on the Add BGP Group Page

Field	Description
Name*	Enter a name for the BGP peer group of the device. The name can contain alphanumeric characters and some special characters [hyphen (-), underscore (_), period (.), and colon (:)] and cannot exceed 64 characters.

Table 54: Fields on the Add BGP Group Page (*Continued*)

Field	Description
Type	Select a type of BGP peer group for the device: • Internal (IBGP) Peer • External (EBGP) Peer
Peer AS	Enter the AS number of the device's BGP peer groups. The value can range from 1 to 4,199,999,999.
BGP Link State	
Originator	Click to enable or disable (default) the BGP peer group as the source for BGP-LS information. If you enable this option, the devices in this group provide the BGP link state information to Routing Director.
Neighbors > IPv4	
Device Labels	Select one or more labels of devices that belong to the BGP peer group. All devices that share the label you enter here become part of the peer group. NOTE: For specifying a single device as a BGP neighbor, provide either the device label or IPv4 address. For specifying multiple devices as a BGP neighbor, use a combination of both device labels and IPv4 and IPv6 addresses. We recommend that you use labels for specifying BGP neighbors as one label can represent multiple devices.
Addresses	Enter the IPv4 addresses (in dotted decimal notation) of the devices that you want to add in the BGP peer group. For example, 10.2.3.4. NOTE: For specifying a single device as a BGP neighbor, you can provide either the device label or IPv4 address. For specifying multiple devices as a BGP neighbor, you can use a combination of both device labels and IPv4 addresses.

Table 54: Fields on the Add BGP Group Page (*Continued*)

Field	Description
Address Family	Select one or more IP address families from the drop-down list that a device can support for BGP sessions with peers.
Neighbors > IPv6	
Device Labels	Select one or more labels of devices that belong to the BGP peer group. All devices that share the label you enter here become part of the peer group. NOTE: For specifying a single device as a BGP neighbor, provide either the device label or IPv4 address. For specifying multiple devices as a BGP neighbor, use a combination of both device labels and IPv6 addresses. We recommend that you use labels for specifying BGP neighbors as one label can represent multiple devices.
Addresses	Enter the IPv6 addresses (in dotted decimal notation) of the devices that you want to add in the BGP peer group. For example, 2001:db8::ff00:42:8329 NOTE: For specifying a single device as a BGP neighbor, you can provide either the device label or IPv6 address. For specifying multiple devices as a BGP neighbor, you can use a combination of both device labels and IPv6 addresses.
Address Family	Select one or more IP address families from the drop-down list that a device can support for BGP sessions with peers.
Route Reflector	

Table 54: Fields on the Add BGP Group Page *(Continued)*

Field	Description
Cluster Name	Select one or more BGP cluster IDs to which you want to assign the devices from the BGP peer group. Click the Manage Clusters link to add, modify, or delete BGP clusters. To add a BGP cluster: 1. Click Manage Clusters. The BGP Route Reflector Clusters page appears. 2. Click the add (+) icon. The Name and Cluster Identifier fields are enabled. 3. Enter a name for the BGP cluster in the Name field. The name can contain alphanumeric characters and some special characters [hyphen (-), underscore (_), period (.), and colon (:)] and cannot exceed 64 characters. 4. Enter an IP address for the BGP cluster in the Cluster Identifier field. Do not enter a value for the cluster ID if you want Routing Director to automatically assign the cluster ID. For automatic configuration of cluster IDs, you must have BGP cluster ID resource pools uploaded to Routing Director. Otherwise, the cluster IDs are not assigned to the BGP clusters, and the device onboarding fails.

Table 55: Fields on the Add Tunnel Page

Field	Description
Name*	Enter a name for the tunnel. The name can contain alphanumeric characters and some special characters [hyphen (-), underscore (_), period (.), and colon (:)] and cannot exceed 64 characters.

Table 55: Fields on the Add Tunnel Page (*Continued*)

Field	Description
Protection	Select the type of protection you want to configure for the tunnel: • none: The tunnel does not have any protection. • link: The links in the tunnel are protected. • node-link: Both the devices and the links in the tunnel are protected. • detour: The tunnel is protected by a secondary tunnel.
Destination	
Device Labels	Select the labels of the devices where you want the tunnel to end. NOTE: You need to provide either the device label or IPv4 address for the tunnel destination. We recommend that you use labels to specify devices for tunnel destination.
Addresses	Enter the IP addresses of the devices where you want the tunnel to end. NOTE: You need to provide either the device label or IPv4 or IPv6 address for the tunnel destination.
Bandwidth	
Bandwidth	Click to enable (default) or disable the automatic configuration (static configuration) of the tunnel bandwidth. If you disable auto configuration (static), specify the tunnel bandwidth in Kbps, Mbps, or Gbps. For example, 5 Mbps.

RELATED DOCUMENTATION

Add Network Resource Pools and Profiles (Day -2 Activities) | 177

Device Onboarding Overview | 171

Add an Interface Profile

An interface profile defines interface-specific configuration elements that are added to the device during onboarding, including the interface's IP address, whether the interface will be used for management or Internet connectivity, or whether the interface will be running OSPF, IS-IS, LDP, or RSVP protocols.

We recommend that you create some interface profiles with configurations that can be applied to all the interfaces that you would add in a network implementation plan and some profiles with interface-specific configurations.

Before you create interface profiles, ensure that you have the required IPv4 address resource pools configured in Routing Director. See ["Add Network Resource Pools and Profiles \(Day -2 Activities\)" on page 177](#) for details.

Routing Director uses the resource pools to assign IP addresses and BGP cluster IDs to the devices.

To add an interface profile:

1. Navigate to **Inventory > Devices > Device and Interface Profiles**.
The Device and Interface Profiles page appears.
2. Click **Add > Interface Profile** to create an interface profile.
The Create Interface Profile page appears.
3. Enter values by referring to [Table 56 on page 222](#).
4. Click **Save** to save the profile.
You can view the profile listed on the Device and Interface Profiles page.

Table 56: Fields on the Create Interface Profile Page

Field	Description
General	
Upload JSON File	Click Browse to upload a pre-created interface profile in the JSON file format. The values in the pre-created interface profile are automatically populated in the Create Interface Profile page. Click the Download this form into JSON file link to download and save the profile in its current state (for example, to save the current configured values for later reference or for maintaining a record).

Table 56: Fields on the Create Interface Profile Page *(Continued)*

Field	Description
Profile Name*	Enter a name for the interface profile. The name can contain alphanumeric characters and some special characters [hyphen (-), underscore (_), period (.), and colon (:)] and cannot exceed 64 characters.
Plan Name	Enter a name for the network implementation plan in which you want to use this profile. You can use the interface profile only in the network implementation plan that you enter here. The name can contain alphanumeric characters and some special characters [hyphen (-) and period (.)] and cannot exceed 64 characters.
Configuration Templates	Select one or more configuration templates that you want to include in the profile. Use configuration templates to configure parameters that you cannot define in the profile or plan. All the configuration templates present in the organization are listed here. When you use a configuration template to configure an interface using NETCONF, you can parameterize only the Interface name in the configuration template. The interface name should be parameterized as <code>{{ interface.name.split('.')[0] }}</code>. See "Add a Configuration Template" on page 379 for details. NOTE: The configurations in the configuration template can be applied only during onboarding. Any changes to the configurations made after the device is onboarded are not applied to the device.
Management	Click to enable or disable (default) the use of an interface as a management interface. If you enable this option, the interface to which you assign this profile is configured as a management interface.

Table 56: Fields on the Create Interface Profile Page *(Continued)*

Field	Description
Internet Connected	Click to enable or disable (default) connectivity tests (by Active Assurance) on an interface. If you enable the Internet Connected option and add the profile as the default interface profile in the network implementation plan, Routing Director initiates connectivity tests from all the ports you configure for all the devices in the network implementation plan. See "Device Connectivity Data and Tests Results" on page 343 for more information. In the network implementation plan, you can also assign the interface profile to specific interfaces (ports).
IPv4 Address	Click to enable or disable (default) the automatic assignment of the IPv4 address for an interface. If you enable this option, Routing Director assigns an IPv4 address to an interface from the resource pool configured in it. For automatic configuration of an IPv4 address, you must have uploaded IPv4 address resource pools to Routing Director. Otherwise, the IP address is not assigned to the device and the device onboarding fails. See "Add Network Resource Pools" on page 676. If you disable this option, you must assign an IPv4 address for the interface in the network implementation plan. See "Add a Network Implementation Plan" on page 241.
IPv6 Address	Click to enable or disable (default) the automatic assignment of the IPv6 address for an interface. If you enable this option, Routing Director assigns an IPv6 address to an interface from the resource pool configured in it. For automatic configuration of an IPv6 address, you must have uploaded IPv6 address resource pools to Routing Director. Otherwise, the IP address is not assigned to the device and the device onboarding fails. See "Add Network Resource Pools" on page 676. If you disable this option, you must assign an IPv6 address for the interface in the network implementation plan. See "Add a Network Implementation Plan" on page 241.
Routing Protocols	

Table 56: Fields on the Create Interface Profile Page *(Continued)*

Field	Description
OSPFv2	Click to enable or disable (default) OSPF configuration on an IPv4 interface. If you enable OSPF configuration, you can configure the Area ID, Metric, and OSPF MTU for the interface.
OSPFv3	Click to enable or disable (default) OSPF configuration on an IPv6 interface. If you enable OSPF configuration, you can configure the Area ID, Metric, and OSPF MTU for the interface.
Area Id	Enter the OSPF area ID for an interface. For example, 0.0.0.1.
Metric	Enter the OSPF metric for the interface. The OSPF protocol uses the cost metric to determine the best path to a destination. Range: 1 through 65,535
OSPF MTU	Enter the maximum transmission unit (MTU) over the OSPF link configured on the interface. Range: 128 through 65,535 bytes
ISIS	Click to enable or disable (default) IS-IS configuration on an interface. If you enable IS-IS, you can configure the IS-IS level, and metric for the interface.
Level	Select the IS-IS level: • IS-IS Level 1 • IS-IS Level 2 • IS-IS Levels 1 and 2
Metric	Enter the IS-IS metric for the interface. The IS-IS protocol uses the cost metric to determine the best path to a destination. Range: 1 through 16,777,215

Table 56: Fields on the Create Interface Profile Page *(Continued)*

Field	Description
LDP	Click to enable or disable (default) LDP configuration on an interface. If you enable LDP, you can enable or disable LDP synchronization for an interface.
LDP Synchronization	Click to enable or disable (default) synchronizing LDP with the underlying IS-IS or OSPF protocol to ensure that LSPs are fully established on an IGP path before forwarding traffic through the LSPs. If LDP is not synchronized with the underlying IS-IS or OSPF protocol, packets might be dropped.
RSVP	Click to enable or disable (default) RSVP configuration on an interface. If you enable RSVP, you can configure link protection for the interface. You must configure this option if you enable traffic engineering (TE) in the device profile that you applied to a device and apply this profile on an interface on the same device.
Link Protection	Click to enable or disable (default) link protection for a tunnel. You must enable link protection if you configure tunnels in the device profile.

RELATED DOCUMENTATION

[Device Connectivity Data and Tests Results | 343](#)

[Device Onboarding Workflow | 173](#)

Add a Port Profile

SUMMARY

Use this topic to add a port profile to Routing Director.

To add a port profile to Routing Director:

1. Navigate to **Inventory > Device and Interface Profiles**.
The Device and Interface Profiles page appears.
2. Click **Add > Port Profile** to create a port profile.
The Create Port Profile page appears.
3. Enter values by referring to [Table 57 on page 227](#).
4. Click **Save** to save the profile.
You can view the profile listed on the Device and Interface Profiles page.

Table 57: Fields on the Create Port Profiles Page

Field	Description
Plan Name	Enter the name of the network implementation plan in which the port profile would be used.
Name	Enter a name for the port profile. A port profile name can contain alphanumeric characters, space, hyphen (-), and underscore (_).
MTU	Enter the maximum transmission unit (MTU) of an aggregated Ethernet interface, in bytes. MTU range depends on the device. Refer to the device documentation for exact range. Range: 128 through 65,535
Tagging	Select a tagging option for an aggregated Ethernet interface: • Tagged—You can assign VLAN tags to the interface. • Untagged—You cannot assign VLAN tags to the interface.

Edit and Delete a Label or Profile

IN THIS SECTION

- [Edit a Label or Profile | 228](#)

Use this topic to edit and delete a label, device profile, or an interface profile.

Edit a Label or Profile



NOTE:

- In a label, the name of the associated plan is not editable.
- In a profile, the names of the profile and the plan to which the profile is assigned are not editable.
- If you edit a profile, you must publish the plan in which the profile is added. Publishing a plan pushes the changes to the devices with which you have associated the profiles. See ["Publish a Network Implementation Plan" on page 258](#).

To edit a label, device profile, or an interface profile:

1. Navigate to **Inventory > Devices > Device and Interface Profiles.**

The Device and Interface Profiles page appears.

2. Select:

- A label and click **Edit** (pencil) icon to edit the label.

The Edit Labels: *Label-Name* page appears.

- A profile and click **Edit** (pencil) icon to edit the profile.

The Edit *Profile-Type: Profile-Name* page appears.

3. Edit the labels and profiles as described in:

- Label, see ["Add Labels" on page 207](#).
- Device profile, see [Fields on the Create Device Profile Page on page 208](#).
- Interface profile, see [Fields on the Create Interface Profile Page on page 222](#).

4. Click **OK to save the label or profile.**

The changes are reflected on the Device and Interface Profiles page.

Delete a Label or a Profile

Before you delete a profile, ensure that the profile is not used in any network implementation plan. To delete a profile from a network implementation plan, edit the plan. See "[Edit a Network Implementation Plan](#)" on page 260.

To delete a label, device profile, or an interface profile:

1. Navigate to **Inventory > Devices > Device and Interface Profiles**.

The Device and Interface Profiles page appears.

2. Select a label or profile and click **Delete** (Trash Can) icon.

A confirmation page appears.

3. Click **Yes** to delete the label or profile.

The label or profile is deleted and removed from the Device and Interface Profiles page.

SEE ALSO

[Device and Interface Profiles Overview](#) | 202

[Network Implementation Plan Overview](#) | 230

Plan Device Onboarding

IN THIS CHAPTER

- [Network Implementation Plan Overview | 230](#)
- [About the Network Implementation Plan Page | 234](#)
- [Add a Network Implementation Plan | 241](#)
- [Publish a Network Implementation Plan | 258](#)
- [Offboard a Network Implementation Plan | 259](#)
- [Edit a Network Implementation Plan | 260](#)
- [Update the Network Implementation Plan After Routing Director Upgrade | 261](#)

Network Implementation Plan Overview

Routing Director uses a network implementation plan to commit configurations on the device during device onboarding, and update configurations after the device is onboarded. For example, if a plan has an RSVP LSP configured from a device to all the provider edge (PE) devices, an LSP is configured from the device to all the PE devices that are currently present in the network and also, to any PE device that might be added to the network after the device is onboarded.

Before you (a network planner) onboard a device, you must create a network implementation plan to define the device configurations to be committed, and health, connectivity, and compliance [with Center for Internet Security (CIS)] checks to be performed on the device.

Network implementation plans define which device profiles, interface profiles, and port profiles should be applied to a device or a group of devices during onboarding. The profiles define which interfaces to configure, with protocols to enable, with IP addresses to assign, and so on.

Routing Director maps the serial number of a device during onboarding with the serial number that is included in a particular network implementation plan and the corresponding device and interface profiles are then applied to the device.

A single network implementation plan can include one or more devices. If you want to onboard multiple devices, you can add all of them to a single implementation plan, and within the plan, reference a

different device profile for each device or specify a default device profile that will be applied to all the devices.




NOTE: You can onboard a maximum of three devices in parallel, at a time.

Additionally, the implementation plan allows the user to provide any required information to build the configuration. For example, for a given interface, you can be referencing an interface profile that does not have automatic address assignment enabled. In the Network Implementation Plan, you will provide the IP address that you want to configure on that interface while creating the network implementation plan.

Routing Director provides Add Network Implementation page in a wizard format that guides you to create the plan. [Figure on page 231](#) shows a screenshot of the Add Network Implementation Plan page designed like a wizard with sections for configuring general information and adding devices and links.

Figure 18: Add Network Implementation Page



To create a plan, navigate to **Inventory > Device Onboarding > Network Implementation Plan**.

In the plan, you:

- Add one or more devices that you want to associate with the plan.
- Assign one or more device profiles and interface profiles to the devices. You can also define default device or interface profiles that would be assigned to all the devices and their interfaces in the network implementation plan. See ["Device and Interface Profiles Overview" on page 202](#) for adding a device or interface profile.

Assigning device and interface profiles to devices is optional. You assign device and interface profiles only if you want to commit IP addresses, protocol configurations, and placement resources on the

device during device onboarding. If you do not assign profiles, you can directly assign the IP addresses and placement resources to the device in the network implementation plan and commit other configurations through configuration templates after the device is onboarded.

- Enter instructions on the type of pluggables to insert and cables to use for connecting to the device ports.

A field technician can view these instructions on the field technician UI while installing the device. So, we recommend that you use terminologies with which a field technician is familiar. See ["Install and Onboard the Device \(Day 0 Activities\)" on page 179](#).


- Add the number of hardware elements (pluggables, memory, PSU, and fans) in the device for collecting health data.
- Configure links from the device to neighboring devices. You can configure links only between devices in the same network implementation plan.



NOTE: You cannot onboard multiple devices in an Implementation plan at the same time.

[Figure on page 233](#) shows an example of a network implementation plan ISP1 created for onboarding an ACX7024 device with serial number 12345.

Figure 19: Example of a Network Implementation Plan



To onboard the ACX7024 device, Routing Director:

- Matches the serial number mentioned in the network implementation plan with the serial number of the ACX7024 device.
- Tags the ACX 7024 device with the label **PE-Router** and configures:
 - PE-1 as the hostname
 - 192.168.1.1 as the loopback interface IPv4 address
- Configures iBGP on the ACX 7024 device based on the PE-Devices device profile and peers the device with other devices with the PE-Router label.

- Configures Interface ge-0/0/0.0 to run OSPF with metric 25 and MTU 9200 based on the CE-FACING interface profile referenced for this interface.
- Configures Interface ge-0/0/1.0 to run IS-IS level 2 with metric 20 and MTU 9200 based on the CORE-FACING interface profile referenced for this interface.

Benefits

- A network architect can provide instructions on the type of pluggables and cables to be used for a port, to a field technician. This helps the field technician to install the correct pluggables and connect correct cables to ports.
- By using a network implementation plan, you can define the configuration for multiple devices once and commit them when the devices are onboarded. To modify the committed configurations on the devices later, you can change the configuration in the plan and push the changes to the devices.
- When you use a plan for onboarding a device, Routing Director executes the health and connectivity checks during device onboarding. The health and connectivity checks during onboarding help you to ensure that the device will function without issues after the device is onboarded and is ready for production soon after onboarding.
- When you use a plan to onboard a device, based on the configurations in the plan, playbooks for collecting metrics are enabled automatically. You do not have to separately configure monitoring. For example, if you enable BGP, Routing Director collects metrics for BGP and displays the data on the Routing Director UI.
- By defining the links between devices in the plan, the links are configured on the devices while the devices are onboarded, enabling quick deployment of your network.

RELATED DOCUMENTATION

[Prepare for Device Onboarding \(Day -1 Activities\) | 178](#)

[Device Onboarding Overview | 171](#)

About the Network Implementation Plan Page

IN THIS SECTION

- [Tasks You Can Perform | 235](#)

To access this page, click **Inventory > Device Onboarding > Network Implementation plan**.

A network implementation plan includes:

- One or more devices to which you assign one or more device profiles and interface profiles. The configurations in the device profiles and interface profiles assigned to the devices are committed on the devices during device onboarding.
- Instructions for installing pluggables in the ports and connecting cables to the device.
- Number of chassis components (fans, power supply modules, pluggables, and line cards) for collecting hardware health data.
- Configuration for links between devices in the plan.

Tasks You Can Perform

You can perform the following tasks from this page:

- View details of a network implementation plan.

To view details of a network implementation plan, select the network implementation plan and click **More > Detail**. Alternatively, hover over the plan name and click the **Details** icon that appears.

The *Network-Implementation-Plan-Name* pane appears on the right of the page displaying information such as the number of devices to which the plan is applied, status of the plan, description and so on. See [Table 58 on page 238](#).

In addition to the fields listed in [Table 58 on page 238](#), you can see the following details on the Network Implementation Plan page:

- Design Version—Design version of the network implementation plan.
- Devices—Devices added to the network implementation plan.
- Device Count—Number of devices included in the network implementation plan.
- Tags—Tags assigned to the network implementation plan.

If there are more than one tag assigned to the plan, an integer is present next to the displayed tag. For example, +2 in the Tags column indicates that there are two more tags in addition to what is currently listed under Tags. You can hover over the integer to view the additional tags assigned to the plan.

- Add a network implementation plan. See ["Add a Network Implementation Plan" on page 241](#).
- Publish a network implementation plan. See ["Publish a Network Implementation Plan" on page 258](#).
- Trigger Onboarding

Use this option to commit the configuration in a network implementation plan on to a device. The device must be onboarded and connected to Routing Director. To commit the network implementation plan configuration onto a device:

1. Click **More > Trigger onboarding**.

The Trigger onboarding page appears.

2. Select the device on which you want to commit the configurations in the network implementation plan and click **OK**.

A confirmation message that the workflow to onboard the device has started appears.

After the configurations are committed successfully, the Order Status changes to **Success** or **Success with Warnings**.

- Resume the Onboarding workflow.

You can use the Resume Onboarding option to resume onboarding a device when the onboarding of a device fails for any reason. Before resuming onboarding, ensure that you have resolved the reason for the onboarding to fail.

To resume the onboarding workflow:

1. Click **More > Resume Onboarding**.

The Resume Onboarding page appears listing the devices included in the plan that have failed onboarding.

2. Select a device from the list and click **OK**.

The onboarding workflow restarts from where it failed and a message indicating that the onboarding workflow has resumed appears.

3. Monitor the progress of the onboarding on the Put Devices into Service page (**Inventory > Device Onboarding > Onboarding Dashboard**).

- Offboard a network implementation plan. See ["Offboard a Network Implementation Plan" on page 259](#).
- Force sync a device.

To use the Force Sync option, click **More > Force Sync**.

The **Force Sync** option reverts a device configuration in the paragon-service-orchestration configuration group to the configuration defined by the network implementation plan that was used to onboard the device. Use this option when you have made changes to the device configuration and want to revert to the configuration defined in the network implementation plan.

- View the history of a network implementation plan.

To view the change history of a network implementation plan, click **More > Order History**. The Order History: *Plan-Name* pane displays the operation (create, modify), status, device count, the user who created or edited the plan and the version of the plan.

Click **View Content** to view the plan in JSON format.

- Export a network implementation plan.

To export a network implementation plan, select the plan and click **More > Export**. The network implementation plan is exported in the JSON format.

The exported file contains all the tasks executed to onboard the devices in the plan. You can export a network implementation plan and use it to:

- View the progress of the onboarding workflow.
- Troubleshoot issues in onboarding.
- Add a tag—A tag is a label in the key:value format that you add to a network implementation plan so that you can group plans. For example, you can assign the tag type:PE to all plans associated with devices deployed as provider edge devices.

To add one or more tags to a device:

1. Click **More > Add Tags**.

The Manage Tags page appears.

2. In the Tags field, select a tag to be applied from the drop-down list.

Alternatively, you can also create a new tag by clicking the Create New Tag link at the bottom of the page and apply the new tag. If you are creating a new tag, ensure that:

- The key starts with a letter and is followed by letters, numbers, periods (.), hyphens (-), or underscores (_), and is between 1 and 200 characters long.
- The value includes alphanumeric characters, periods (.), hyphens (-), and underscores (_), and is between 1 and 200 characters long.

You can view all the tags applied to the network implementation plan in the Tags column.

- Upgrade the infrastructure service design used by a network implementation plan.

To upgrade the infrastructure service design used by an implementation plan:

1. Select the plan and click **More > Upgrade Service Design**.

The Upgrade Service Design page appears.

2. On the Upgrade Service Design dialog box, click the drop-down list to select the service design version to which you want to upgrade the network implementation plan and click **Yes**.

A message indicating that service design upgrade is in process is displayed. After successful upgrade, the Service Design column displays the new version of the infrastructure service design used by the network implementation plan.

- Search by using keywords—Click the search icon (magnifying glass), enter the search term in the text box, and press Enter. The search results are displayed on the same page.
- Filter the data displayed in the table—Click the filter icon (funnel) and select whether you want to show or hide advanced filters. You can then add or remove filter criteria, save criteria as a filter, apply or clear filters, and so on. The filtered results are displayed on the same page.
- Show or hide columns in the table or reset page preferences, using the vertical ellipsis menu.
- Sort, resize, or re-arrange columns in a table (grid).
- Refresh the Network implementation Plan table and set auto refresh frequency—A Refresh icon is provided at the bottom left-corner of the Network Implementation Plan table. Click the **Refresh** icon to view the most recent updates made to the network implementation plans.

Click the **ellipsis** next to the Refresh icon to set the auto refresh frequency. You can either turn off the auto refresh of the table or set the auto refresh frequency to 30 seconds, 1 min, 3 mins, and 5 mins.

- Set the number of rows to be displayed in the table—Set the number of rows (15, 30, 50, or 100) to be displayed in the Network Implementation Plan table in the **Display** drop-down list provided at the bottom right-corner of the table.

If there are multiple pages of rows, use the arrow heads to traverse through the pages or enter a page number to view the entries listed on that page.

Table 58: Fields on the Network Implementation Plan Page

Field	Description
Properties	
Name	Name of the network implementation plan.
Description	Short description of the network implementation plan.
Sites	Name of the site where the devices are installed.

Table 58: Fields on the Network Implementation Plan Page *(Continued)*

Field	Description
Order State	The execution state of the network implementation plan. A network implementation plan can have the following execution states: • Processing—The network implementation plan is in the process of being executed. • Success—The network implementation plan is successfully executed. • Success with warnings—The network implementation plan is successfully executed but certain tasks in the onboarding workflow have failed. • Not Applicable—The onboarding workflow related to the network implementation plan failed to execute.
Status	Status of the network implementation plan. A network implementation plan can have one of the following provisioning states: • Feasibility Checked—The network implementation plan is created, • Designed—Possible resource allocation options are assigned for the network implementation plan, but the resources are not yet reserved or allocated. • Reserved—Resource allocation for the network implementation plan is complete. • Provisioning—Onboarding workflow related to the network implementation plan is being executed with the allocated resources. • Inactive—The network implementation plan is deactivated. • Active—Device configurations in the network implementation plan were successfully committed on the devices. • Blocked—Device configurations could not be applied on one or more devices in the network implementation plan. • Terminated—The network implementation plan is deleted.
Status Message	Additional information about the status of the network implementation plan. The status is Uploaded after the network implementation plan is saved either after creating or modifying.

Table 58: Fields on the Network Implementation Plan Page (*Continued*)

Field	Description
Device Profiles	Device profiles included in the plan.
Interface Profiles	Interface profiles included in the plan.
Port Profiles	Port profiles included in the plan.
Last Modified	Date and time in the Month Day, Year HH:MM:SS format when the network implementation plan was last modified.
Last Modified By	User who last modified the network implementation plan.
Devices	View the devices included in the plan along with the information if the device is onboarded or not. You can: - Click the <i>device-name</i> link to see the result of device health checks on the Overview tab of the <i>Device-Name</i> page (Observability > Troubleshooting Devices > Click a <i>device-name</i>). - Click the More link in a device tile to view the Device Detail - <i>Device-Name</i> page. On the Device Detail - <i>Device-Name</i> page, you can see: - Properties—View the device properties (for example, device hostname, IPv4 loopback address, device model, and so on) - States—View the number of the active assurance tests performed and the health of the device components (such as linecards, fans, power supply modules, optics, and so on).

Table 58: Fields on the Network Implementation Plan Page *(Continued)*

Field	Description
More	You can see: • JSON View—View the network implementation plan in the JSON format. • Order Status—View the components affected and workflow tasks executed when the infrastructure service order is executed. • Placement Resources—View the placement resources (bandwidth, MAC addresses, routes, and integrated Routing and Bridging) configured in the network implementation plan. • Inventory—View the Inventory page (Inventory > Devices > Network Inventory).

RELATED DOCUMENTATION

[About the Device and Interface Profiles Page | 205](#)

[View Results of Automated Device Tests | 268](#)

Add a Network Implementation Plan

You must be a user with the Network Admin or Super User role to add a network implementation plan.

To add a network implementation plan:

1. Navigate to **Inventory > Device Onboarding > Network implementation Plan**.
The Network Implementation Plan page appears.
2. On the Network implementation Plan page, do one of the following:
 - Select the implementation plan that was created automatically after you created the device plan (the name of the plan will be the name you entered in the device profile), and then click **Edit** (Pen) icon.
 - Click **+** (Add) to create a network implementation plan.

If you create a new plan instead, the device profiles that you created before will not be available for selection within the implementation plan.
3. Enter values by referring to [Table 59 on page 242](#).
4. Click **Save** to save the plan.

The plan is listed on the Network Implementation Plan page.

After you create a network implementation plan, the devices included in the plan can be installed and onboarded to Routing Director.

Table 59: Fields on the Add Network Implementation Plan

Field	Description
General	
Upload JSON File	Click Browse to import a pre-created network implementation plan in JSON format. The values in the pre-created plan are automatically populated in the Add Network Implementation Plan page. Click the Download this form into JSON file link to download and save the profile in its current state (for example, when you want to save the current configured values for later reference or for maintaining a record).
Plan Name*	If you are creating a plan, enter a name for the plan. The plan name can contain alphanumeric characters (a-z, A-Z, and 0-9) and some special characters [period (.) and hyphen (-)], and cannot exceed 64 characters. If you are editing a plan, the name is already populated and cannot be edited.
Description	Enter a description for the plan.
Default Device Profile	Select one or more device profiles to be used in the plan. You can view only those device profiles that you associate with the plan while creating the profile. If you are editing an automatically generated implementation plan, the default interface and device profiles are already populated. Configurations in the default device profile are common to all devices and applied to all the devices included in the plan. Alternatively, click the Add new device profile link to create a device profile to be used as the default device profile. See "Add a Device Profile" on page 208.

Table 59: Fields on the Add Network Implementation Plan *(Continued)*

Field	Description
Default Interface Profile	Select one or more interface profiles to be used in the plan. You can view only those interface profiles that you associate with the plan while creating the profile. The configurations in the default interface profile are common to all interfaces and applied to all the interfaces configured in the plan. Alternatively, click the Add new interface profile link to create an interface profile to be used as the default interface profile. See "Add an Interface Profile" on page 222.
Default Port Profiles	Select one or more port profiles to be used in the plan. You can view only those port profiles that you associate with the plan while creating the profile. The configurations in the default port profile are committed on aggregated Ethernet interfaces that are included in the plan. Alternatively, click the Add new port profile link to create a port profile that can be used as the default port profile. See "Add a Port Profile" on page 226.
Devices	Add all devices that you want to configure while the device is onboarded to Routing Director. The plan can also be used to manage the configuration of the devices after the devices are onboarded. You can also edit and delete the devices added to the plan from here. To add devices: - Click the Add (+) icon to add devices to the plan. The Add Devices wizard appears that helps you configure the device, the device's interfaces, and add the chassis components for monitoring health. See Table 60 on page 244 for adding devices. - Click the Import Devices button and select the devices to be added to the network implementation plan from the List of Devices page. NOTE: You can add a device to only one network implementation plan.
Links	Add links between the devices added to the plan. Click the Add (+) to add links to other devices. The Add Link page appears from where you can configure the links. You can also edit and delete the links configured between devices included in the plan from here. See Table 62 on page 257.

Table 59: Fields on the Add Network Implementation Plan *(Continued)*

Field	Description
Summary	Displays a summary of the onboarding plan. Click the Edit link to edit the general information or the links added to the plan.

Table 60: Fields on the Add Device Page

Field	Description
General	
Name*	Enter a name for the device. Routing Director uses this name internally. The name can contain alphanumeric characters and some special characters [hyphen (-) and underscore (_)]. The name cannot exceed 64 characters. Alternatively, enter the hostname of the device and click the Search icon to populate the device values from the network inventory on this page.
Hostname	Enter a hostname for the device. The name can contain alphanumeric characters and some special characters [hyphen (-) and underscore (_)]. The hostname cannot exceed 64 characters. If you do not enter a hostname, Name is used as the hostname.
IPv4 Loopback	Enter an IPv4 loopback address for the device in the dotted decimal notation format. For example, 10.10.10.1. The value that you enter here overrides the value that is configured in a device profile.
IPv6 Loopback	Enter an IPv6 loopback address for the device. For example, 2001:db8:4136:e378:8000:63bf:3fff:fdd2. The value that you enter here overrides the value that is configured in a device profile.
Site*	Select the site where you want to install the device. Alternatively, if you have permissions to add a site, you can view the Add new site link next to the Site drop-down list. Click the link and add a new site on the Create Site page. See "Add Sites" on page 90.

Table 60: Fields on the Add Device Page (*Continued*)

Field	Description
Serial Number	Enter the serial number of the device that you want to associate with the plan. The serial number will be used to map the device to this profile when it is added to the inventory (during adoption which is described later), and the onboarding process is started.
Vendor	Select the vendor of the device.
Model*	Select the model of the device from the drop-down list. For a list of supported devices, see Supported Junos OS Releases, Devices, and Browsers.
Software Image	Select the software image to be installed on the device during onboarding from the drop-down list. All software images that are uploaded to Routing Director are listed here.
Device Profiles	Select one or more device profiles to be applied to the device from the drop-down list. The configurations in the device profiles are committed on the device in the order in which the profiles are added to the plan. Configurations present in both the default device profile and the specific profiles that you enter here are committed on the device. However, for configurations that are present in both the specific device profiles and the default device profile, the values in the specific device profiles override the configuration in the default device profile.

Table 60: Fields on the Add Device Page (*Continued*)

Field	Description
Installation Duration	Enter the time duration in minutes that the field technician can take to install the device, add pluggables, and connect cables. Routing Director checks the health of the device as soon as it pushes the configurations on the device. However, If the device installation is not complete by then, Routing Director notifies the device as unhealthy and executes the next task. To prevent reporting the device as unhealthy before installation is complete, Routing Director retries the device health checks every minute. The device health is checked until the device is detected as healthy or for the duration that you enter here, whichever is earlier. Default: 20 minutes Range: 0 through 20
Instructions	Enter instructions or any additional information that you want to provide to a field technician who is onboarding the device. The field technician can view the instructions that you enter here.
Placement Resources Enable or disable (default) allocating resources for provisioning services on the device.	
Bandwidth*	Enter the bandwidth (in Mbps) to be reserved on the device for service provisioning. NOTE: If you are configuring a PE node, you must configure bandwidth for the PE node to appear as a resource for a service.
MAC Address	Enter the maximum number of MAC addresses to be assigned to the device for service provisioning.
Routes	Enter the maximum number of routes addresses to be assigned to the device for service provisioning.
IRBs Enable or disable (default) integrated routing and bridging (IRB) interfaces on a PE device.	

Table 60: Fields on the Add Device Page (*Continued*)

Field	Description
Start*	Enter the starting IRB for the range of IRB IDs to which the PE interface connects. Enter a value in the range 1 through 65535.
Size*	Enter the number or range of IRB IDs available for a PE device. You can define a range using a starting ID and the number of IRBs required. Enter a value in the range 1 through 65535.

Table 60: Fields on the Add Device Page (*Continued*)

Field	Description
Physical Ports	Displays the device's chassis view. You can perform the following tasks from here: - View the ports (interfaces) on the device chassis in the displayed chassis image. - Add and define the configuration of all the logical ports that you want to configure during device onboarding. To add a logical port: - Click the Add Default Interfaces to populate the interfaces in the device. - Click a port that you want to configure and click the Edit icon (pen). Alternatively, click the Add (+) icon to specify and configure a port. See "Configure Port" on page 248. NOTE: In this release, Interfaces are populated for the ACX7024 devices only. - Add and define the configuration of all aggregated Ethernet interfaces that you want to configure during device onboarding. To add and define an aggregated Ethernet interface, click Add icon to specify and configure the port. See Table 61 on page 253. NOTE: To migrate a service from an earlier release of Routing Director to a later release, services using ports with aggregated Ethernet interfaces must add the port in the network implementation plan with its access parameters configured. The access parameters must not have VLAN configuration. - Enter the number of LAGs that can be configured on the device (in the Count field under Link Aggregation Group (LAG) Settings). The maximum number of LAGs supported on a device depends on the device type. - Edit, or delete the ports that you added. - Zoom in and zoom out the chassis view.
Configure Port	

Table 60: Fields on the Add Device Page (*Continued*)

Field	Description
Interface Name*	Enter a name for the interface as follows: • Use the type-fpc/pic/port.logical format to enter the name with a logical unit. For example, et-0/0/0.0. • Use fpc/pic/port[:channel].logical format to enter the name for a channelized interface. For example, t3-0/0/0:1.0. • Use the management interface name of the device for management interfaces. For example, re0:mgmt-0.0. • Use the aggregated Ethernet interface name for aggregated interfaces. For example, ae0, ae1, and ae2. NOTE: Unit 0 is the only logical unit supported for an interface. For ports in the ACX7204 devices, this field is auto-populated if you select a port to be configured in the Physical Ports tab.
Description	Enter a description for the interface.
IPv4 Address/Subnet Mask	Enter the IPv4 address (in dotted decimal notation) with the subnet mask for the interface. For example, 10.10.10.10/24. If you have disabled automatic IP address assignment in the interface profiles assigned to the interface, you can assign the IPv4 address for the interface here.
IPv6 Address/Subnet Mask	Enter the IPv6 address with the subnet mask for the interface. For example, 2001:db8:4136:e378:8000:63bf:3fff:ddd2/48 If you have disabled automatic IP address assignment in the interface profiles assigned to the interface, you can assign the IPv6 address for the interface here.

Table 60: Fields on the Add Device Page (*Continued*)

Field	Description
Interface Profiles	Select one or more interface profiles to be applied to the interface from the drop-down list. The configurations in the device profiles are committed on the device in the order in which the profiles are added to the plan. Configurations present in both the default interface profile and the specific profiles that you enter here are committed on the device. However, for configurations that are present in both the specific interface profiles and the default interface profile, the values in the specific interface profiles override the configuration in the default interface profile.
Pluggable	Enter the type of pluggable to use in the port; for example, QOD-400G-FR4. A field technician can view the information that you enter here during device onboarding.
Cabling Instructions	Enter instructions to connect cables to the interfaces. A field technician can view and use this information that you enter here to connect cables to interfaces during device onboarding. We recommend that you use specific instructions that are known to the field technician, such as a reference number of the cable.
Access Parameters Click the add (+) icon above the Access Parameters table. The Add Access Parameter page appears.	
Name*	Enter a reference name for the CE facing interface.
Tag Type	Select dot1q, qinq, or untagged as the VLAN tag from the Tag Type drop-down list. The default tag type is untagged.

Table 60: Fields on the Add Device Page (*Continued*)

Field	Description
dot1q Tag Type	A Dot1q interface transmits and receives Ethernet frames with a single VLAN tag. When you select dot1q from the Tag Type drop-down, the VLANs section appears. Enter VLAN details by referring to the following list: • Start*—Enter the starting VLAN ID for the range of VLANs from the CE device that the PE interface supports. Enter a value in the range 1 through 4094. For example, if 200 is the starting VLAN ID, the first VLAN supported by the PE interface has a VLAN ID of 200. • Size*—Enter the range of VLANs from the CE device that the PE interface supports. Enter a value in the range 1 through 4094.

Table 60: Fields on the Add Device Page *(Continued)*

Field	Description
qinq Tag Type	A Q-in-Q interface transmits and receives frames with double VLAN tags. When you select qinq from the Tag Type drop-down, the Inner VLANs and Outer VLANs sections appear. Enter Inner VLANs and Outer VLANs details by referring to the following list: Enter Inner VLANs details: • Start*—Enter the starting VLAN ID for the range of inner VLANs. Enter a value in the range 1 through 4094. For example, if 200 is the starting inner VLAN ID, the first inner VLAN supported by PE device has a VLAN ID of 200. • Size*—Enter the range for inner VLANs. Enter a value in the range 1 through 4094. Enter Outer VLANs details: • Start*—Enter the starting VLAN ID for the range of outer VLANs. Enter a value in the range 1 through 4094. For example, if 200 is the starting outer VLAN ID, the first outer VLAN supported by PE device has a VLAN ID of 200. • Size*—Enter the range for outer VLANs. Enter a value in the range 1 through 4094.
untagged Tag Type	Select untagged tag type to assign no VLANs to the interface. An untagged interface receives and sends Ethernet frames without VLAN tags.
Chassis	
PSMs	Enter the number of power supply modules (PSMs) in the device. This information is used for collecting analytics and health data from the PSMs.

Table 60: Fields on the Add Device Page (*Continued*)

Field	Description
Fans	Enter the number of fans in the device. This information is used for collecting analytics and health data from the fans.
Linecards	Enter the number of line cards in the device. This information is used for collecting analytics and health data from the line cards.
Pluggables	Enter the number of pluggables in the device. This information is used for collecting analytics and health data from the pluggables.

Table 61: Fields on the Configure Port page (for configuring aggregated Ethernet interfaces)

Field	Description
Name	Enter a name for the aggregated Ethernet interface. For example, ae1, ae3.
Description	Enter a description of the interface. For example, Link to PE router.
MTU	Enter the maximum transmission unit (MTU) for the aggregated Ethernet interface in bytes. The MTU range depends on the device. Refer to the device documentation for the exact range. Range: 128 through 65,535
Tagging	Select a tagging option for the aggregated Ethernet interface: - Tagged – You can configure VLANs on the interface. Define the type of tagging for the port under Access Parameters. - Untagged – You cannot configure VLAN on the interface.

Table 61: Fields on the Configure Port page (for configuring aggregated Ethernet interfaces)
(Continued)

Field	Description
Port Profiles	Select one or more port profiles to be assigned to the aggregated Ethernet interface. Click Add new port profile to configure a new port profile and assign it to the interface. For the same values defined in the port profile and the network implementation plan, the values defined in the plan take precedence over the values in the profile.
LAG Parameters NOTE: Link Aggregation Control Protocol (LACP) is not supported in this release.	
Minimum Links	Enter the minimum number of links that can be configured on the aggregated Ethernet interface. Range: 1 through 64
Member Link	
Name	Enter a name for the physical interfaces (Ethernet interfaces and channelized interfaces) that are a part of the aggregated Ethernet interface, For example, ge-1/0/0, ge-1/0/1. t3-0/0/0:1.2
Access Parameters Add the range of VLANs and the number of VLANs for service provisioning on the aggregated Ethernet interface. If you assign multiple VLANs, you can configure multiple services on the LAG. If you do not configure a VLAN, you can configure a single service on the LAG. To use an existing manually configured aggregated Ethernet interface on a device as an access interface for a service, define the interface in the network implementation plan by including VLAN in the interface's access parameters.	
Name*	Enter a reference name for the CE facing interface.

Table 61: Fields on the Configure Port page (for configuring aggregated Ethernet interfaces)
(Continued)

Field	Description
Tag Type	Select the type of VLAN tagging to be used on the aggregated Ethernet interface: • Untagged • dot1q • qinq If you select dot1q tag type, enter the VLAN starting ID and VLAN size. If you select qinq tag type, enter the VLAN starting ID and VLAN size for outer (service) VLAN and inner (customer) VLAN.
VLANs	
Start*	Enter the starting VLAN ID for the range of VLAN IDs that will be assigned to the interface. For example, a value of 10 indicates that the range of VLAN IDs assigned to the interface starts from 10. Range – 1 through 4094
Size*	Enter the size of the VLAN range that will be assigned to the interface. For example, a size of five with starting VLAN ID as 10 indicates that the VLAN IDs 10 through 14 are assigned to the interface. Range: 1 through 4094
q-in-q	
Inner VLANs	

Table 61: Fields on the Configure Port page (for configuring aggregated Ethernet interfaces)
(Continued)

Field	Description
Start*	Enter the range of inner VLAN IDs (customer VLAN IDs) that will be assigned to the interface. For example, a size of five with starting inner VLAN ID as 15, indicates that the inner VLAN IDs 15 through 19 are assigned to the interface. Range: 1 through 4094
Size*	Enter the starting outer VLAN ID (service VLAN ID) for the range of outer VLAN IDs that will be assigned to the interface. For example, a value of 25 indicates that the range of outer VLAN IDs assigned to this interface starts from 25. Range: 1 through 4094
Outer VLANs	
Start*	Enter the starting outer VLAN ID (service VLAN ID) for the range of outer VLAN IDs that will be assigned to the interface. For example, a value of 25 indicates that the range of outer VLAN IDs assigned to this interface starts from 25. Range: 1 through 4094
Size*	Enter the range of outer VLAN IDs (service VLAN IDs) that will be assigned to the interface. For example, a size of five with starting outer VLAN ID as 25, indicates that the outer VLAN IDs 25 through 29 are assigned to the interface. Range: 1 through 4094

Table 62: Fields on the Add Link Page

Field	Description
Link Name	Enter a name for the link. The name can contain alphanumeric characters and some special characters [hyphen (-), underscore (_), period (.), and colon (:)] and cannot exceed 64 characters. You must enter the link name if you want to configure links between multiple devices in the same subnet.
Device A	
Device*	Select a source device to originate the link.
Site	Displays the site where the device that originates the link is installed.
Interface*	Select the interface on the source device from which the link originates.
Connection Instructions	Enter instructions for the link. For example, the cables to be used to connect the device to the network or another device.
Device Z	
Device	Select the destination device to terminate the link. You need not select a destination device if you want to connect to multiple devices from the same source device and interface.
Site	Displays the site where the destination device that terminates the link is installed.
Interface	Select an interface on the destination device at which the link terminates. You need not select a destination interface if you want to connect to multiple devices from the same source device and interface.
Connection Instructions	Enter instructions for the link. For example, the cables to be used to connect the device to the network or another device.

RELATED DOCUMENTATION

[Add a Device Profile | 208](#)

[Install and Onboard the Device \(Day 0 Activities\) | 179](#)

[Trust and Compliance Overview | 603](#)

[Device Onboarding Workflow | 173](#)

Publish a Network Implementation Plan

You must be a user with the Network Admin or Super User role to publish a network implementation plan.

You publish a plan after you modify:

- Data included in the plan
- Data in any of the profiles included in the plan

When you publish a plan, the changes are pushed to the respective devices that are included in the plan. When you save the plan, the configuration is saved only in the Routing Director database and not pushed to the devices.

When you edit a plan, the status of the plan is Uploaded. After the plan is published, the state changes to Transformed.

To publish a network implementation plan:

1. Click **Inventory > Device Onboarding > Network Implementation Plan**.
The Network Implementation Plan page appears.
2. Select a plan that you want to publish and click **Publish**.
A confirmation message appears indicating that the plan is published. The status of the plan is changed to Transformed.
3. (Optional) Log in to the device and view the configurations to confirm that the published changes are applied.

RELATED DOCUMENTATION

[Network Implementation Plan Overview | 230](#)

[Device Life-Cycle Management Overview | 168](#)

Offboard a Network Implementation Plan

You must be a user with the Network Admin or Super User role to offboard a device.

Use the offboard option when you want to delete a network implementation plan and stop Routing Director from managing devices assigned to the plan. When you offboard a network implementation plan:

1. The configurations applied on the device through the plan are deleted.
2. The plan is deleted from the database and is no longer listed on the Network Implementation Plan page.



NOTE:

- The outbound SSH command committed on the device is not deleted when you offboard a network implementation plan. To delete the outbound SSH configuration, you must release the device. See ["Release a Device" on page 133](#).

You can also offboard a device by deleting the device (on the Devices tab) from the network implementation plan.

- You cannot offboard a plan when the plan is being used to onboard a device.

To offboard a network implementation plan:

1. Navigate to **Inventory > Device Onboarding > Network Implementation Plan**.

The Network Implementation Plan page appears.

2. Select one or more plans that you want to delete or offboard.

3. Click **More > Offboarding**.

A confirmation dialog box appears.

4. (Optional) Clear the following check boxes:

- **Also delete the profiles and labels associated with the plan.**

This option is selected by default and deletes the profiles and labels associated with the network implementation plan. Clear the check box to avoid deleting the profiles and labels associated with the plan.

- **Ignore unreachable devices.**

This option is selected by default and ignores devices included in the plan that are unreachable during offboarding. If you clear this option and if the devices in the plan are unreachable during offboarding, the plan offboarding fails.

5. Click **Yes**.

The offboard process is initiated. After the offboarding is complete, the offboarded plans are no longer listed on the Network Implementation Plan page.

6. (Optional) Delete the devices that were included in the offboarded plan from the Inventory page (**Inventory > Devices > Network Inventory**). See ["Release a Device" on page 133](#). Releasing a device deletes the outbound SSH configuration from the device and the device is not connected with Routing Director.

After you release a device, the device is no longer listed on the Inventory page.

RELATED DOCUMENTATION

[Add a Network Implementation Plan | 241](#)

[Device Life-Cycle Management Overview | 168](#)

Edit a Network Implementation Plan

You must a user with the Network Admin or Super User role to edit and publish a network implementation plan.



NOTE:

- You cannot edit a network implementation plan when a device is being onboarded by using that plan.
- You should not attempt to manually modify the device configuration in the paragon-service-orchestration, jcloud-gnmi-sensors, and jcloud-script configuration groups. The configurations might not work if you modify manually.

To edit a network implementation plan:

1. Navigate to **Inventory > Device Onboarding > Network Implementation Plan**.
The Network Implementation Plan page appears.
2. Select a plan and click **Edit** (pencil) icon.
3. Enter values by referring to [Table 59 on page 242](#).



NOTE: You cannot edit the Plan Name.

4. Click **Save** to save the plan.

You can view the changes on the Network Implementation Plan page.

5. Click **Publish** to publish the plan.

Publishing an edited plan ensures that the edits are applied to the devices included in the plan.

6. (Optional) Click **Export** to view the progress of the updates being made to the devices.

RELATED DOCUMENTATION

[Publish a Network Implementation Plan | 258](#)

[Network Implementation Plan Overview | 230](#)

Update the Network Implementation Plan After Routing Director Upgrade

SUMMARY

After the Routing Director cluster has been upgraded, use the information in this topic to upgrade the service designs, and update the network implementation plan on the upgraded setup.

1. As part of upgrading your Routing Director installation from an earlier release, you must also upgrade the service designs that Routing Director provides for infrastructure, resources (I2-addr, I3-addr, routing, topo, and vpn), and services (L3VPN, EVPN, EVPN-VPWS, and L2 circuit). The upgrade of your Routing Director installation is incomplete without upgrading the service designs.

After you upgrade your Routing Director installation to Release 2.5.0, upgrade the service, resource, and onboarding service designs versions to the following:

- elan-evpn-csm—0.5.54
- eline-l2circuit-nsm—0.5.22
- eline-evpn-vpws-csm—0.0.30
- l2vpn—0.3.41
- l2vpn-nsm—0.2.22
- l3vpn—0.7.73
- l2-addr—0.1.7

- l3-addr—0.2.10
- routing—0.1.11
- topo—0.6.51
- vpn-resources—0.3.7
- infrastructure—0.14.46
- iana—0.1.1
- ietf—0.1.1
- ietf-draft—0.1.1
- juniper—0.1.9

See ["Upgrade Service Design Version for a Service Instance" on page 658](#) to upgrade the service design versions.

2. Update the network implementation plan created in an earlier release to use the infrastructure service design version 0.14.46.

- a. Log in to the Routing Director GUI and obtain the organization ID.

- i. Click **Settings Menu > System Settings** on the banner.

The Organization Settings page appears.

- ii. Click **Copy** under Organization ID field.

- b. Select a plan and upgrade the service design to 0.14.46.

- i. On the Routing Director GUI, click **Inventory > Device Onboarding > Network Implementation Plan**.

The Network Implementation Plan page appears.

- ii. Select the plan that you want to update and click **More > Upgrade Service Design**.

The Upgrade Service Design page appears.

- iii. Select 0.14.46 from the drop-down list and click **Yes**.

A confirmation message appears that the service design of the selected plans is upgraded to the selected service design version.

- c. If there are multiple network implementation plans, repeat steps ["2.b" on page 262](#) for all plans.

- d. (Optional) On the Routing Director GUI, select the plan that you upgraded and edit the new parameters associated with the upgraded service design.
- e. Click **Publish** for the changes in the plan to be pushed to the devices in the plan.

When you click Publish,

- Resources are allocated again for the devices in the plan and
 - The service design version recorded in Placement Resources is matched with the service design version of the plan and the new configurations available with the new service design are pushed to the devices.
- f. Verify that the publish workflow for the plan is executed successfully and the Order Status of the plan is *Success* on the Network Implementation Plan page.

RELATED DOCUMENTATION

[Service Design Overview | 651](#)

[Network Implementation Plan Overview | 230](#)

View Device Onboarding

IN THIS CHAPTER

- [About the Put Devices into Service Page | 264](#)
- [Move a Device to Production | 268](#)
- [View Results of Automated Device Tests | 268](#)
- [Identity and Location Data of a Device | 270](#)
- [Remote Management Data and Test Results | 272](#)
- [Hardware Data and Test Results | 278](#)
- [Interfaces Data and Test Results | 287](#)
- [Software Data and Test Results | 306](#)
- [Configuration Data and Test Results | 308](#)
- [Routing and MPLS Data and Test Results | 310](#)
- [Device Connectivity Data and Tests Results | 343](#)

About the Put Devices into Service Page

IN THIS SECTION

- [Tasks You Can Perform | 265](#)
- [Field Descriptions | 266](#)

To access the Put Devices into Service page, navigate to **Inventory > Device Onboarding > Onboarding Dashboard**.

The Put Devices into Service page is a dashboard in Routing Director. It helps you to view a summary of devices that:

- Are ready to install, which means that all pre-requisites to install the device at a site are complete.
- Need urgent action, which means that the device has some critical issues. Immediate user intervention is needed for resolution.

You can also view a comparison (as a number or percentage) of alerts generated in the current week against those in the past week. Hover over the widget to view the number of critical alerts generated in the current week and in the past week.

- Need action, which means that the device has major issues. Immediate user intervention is not needed for resolution; for example, a disk in a device is not booting.

You can also view a comparison (as a number or percentage) of major alerts generated in the current week against the alerts in the past week. Hover over the widget to view the number of major alerts generated in the current week and in the past week.

You can use the text boxes and drop-down lists below the table headers to filter devices by the respective attribute values. By default, devices not in service yet are displayed. This can be changed by using the **Operational State** column filter.

Some filters additionally display the device count by group. For example, you can click the Sites filter and view the number of devices present at each configured site.

Tasks You Can Perform

The tasks that you can perform on the Put Devices into Service page are:

- Filter devices by using the table headers—Device's hostname, model number, IP address, site, status, OS version, and so on.

Click the search icon (magnifying glass), in the text box or drop-down list under each header. Enter the search term (hostname, model number, IP address, OS version) in the text box or the search value in the drop-down list, and press Enter. You can view the search results on the same page.

In the list of devices filtered, click the *hostname* link to view the device details and test results.

- View the results of automated checks performed on the device during onboarding.

To view the results, click the device's *hostname* link. The *Device-Name* page appears displaying the result of all the tests executed during onboarding. In addition, when the device is in production, you can continue to monitor the device from the *Device-Name* page. For more information, See ["View Results of Automated Device Tests" on page 268](#).

- Put the device into production; See ["Move a Device to Production" on page 268](#).
- Sort, resize, or re-arrange columns in a table (grid).
- Show or hide columns in the table or reset page preferences, using the vertical ellipsis menu.

- Delete a device. See ["Release a Device" on page 133](#).
- Set the number of rows to be displayed in the table—Set the number of rows (15, 30, 50, or 100) to be displayed in the Network Implementation Plan table in the **Display** drop-down list provided at the bottom right-corner of the table.

If there are multiple pages of rows, use the arrow heads to traverse through the pages or enter a page number to view the entries listed on that page.

Field Descriptions

[Fields on the Put Devices into Service Page on page 266](#) describes the fields on the Put Devices into Service page.

Table 63: Fields on the Put Devices into Service Page

Field	Description
Hostname	Hostname assigned to the device. Click the hostname to open the <i>Device-Name</i> page and view the results of the automated tests conducted on the device during onboarding; see "View Results of Automated Device Tests" on page 268. You can continue monitoring the health and performance of the device from the <i>Device-Name</i> page when the device is in production.
Severity	Indicates the severity of the events on the device. The severity of the events are categorized as: • Urgent Action Needed (critical)—Indicates that your action or intervention is needed to resolve one or more issues that are affecting the functioning of the device. • Action Needed (major)—Indicates that a major event has occurred on the device. The functioning of the device is affected in some way, but not drastically. You might need to intervene to resolve the issue, but not immediately. • Being Monitored (minor)—Indicates that a minor event has occurred on the device and that the device is being monitored. • Healthy—Indicates that the device is healthy and that there are no issues with the health and functioning of the device.

Table 63: Fields on the Put Devices into Service Page *(Continued)*

Field	Description
IPv4 Address	IPv4 management address configured on the device.
IPv6 Address	IPv6 management address configured on the device.
Model	Model of the device; for example: ACX7100-48L.
Serial Number	Serial number of the device.
OS Version	Version of the OS installed on the device; for example, 22.2R1.9.
Site	Site (geographical location) where the device is installed.
Type	Function of the device in the network.
Connected	Indicates whether the device is connected (yes) or not connected (no) to Routing Director.
Operational State	Status of the device: • Undefined—The state of the device is not defined by Routing Director. This is the default state. • Onboarding—The device is in the process of being onboarded to Routing Director. • Ready for Service—The device is installed and ready to be moved to production. • In Service—The device is in production. Services are reconfigured on the devices so that the device can route production traffic. • Maintenance—The device is undergoing maintenance; for example, software upgrade. • RMA—The device is getting replaced by another device.

RELATED DOCUMENTATION

Move a Device to Production

You must be a user with Super User or a Network Admin role to move the device from the Ready for Service state into production. You can move a device to production only when the device does not have any critical or major alerts. To move the device into production, navigate to Inventory > Onboarding Dashboard. Select the device and click **Put into Service** present above the devices table. The status of the device changes to In Service indicating that the device can transport traffic.

Alternatively:

1. Click the *hostname* link of the device that you want to put into service on the Put Devices into Service page.
The *Device-Name* page appears.
2. Click **Put into Service Now** to move the device to production.

After you move the device to production, you can monitor the device's performance from the *Device-Name* page.

RELATED DOCUMENTATION

[Service Orchestration Overview | 646](#)

[Troubleshoot Using Alerts and Alarms | 450](#)

View Results of Automated Device Tests

After a device is connected to Routing Director, Routing Director executes a series of automated tests to verify the device onboarding. For example, tests are executed to check the health of the CPU and memory, connectivity with the neighboring devices, the remote management connection between the device and Routing Director, and so on. You can view the results of the tests on the *Device-Name* page.

You can access the *Device-Name* page by clicking the hostname of the device on the Put the Devices Into Services page (**Inventory > Device Onboarding > Onboarding Dashboard**) or Troubleshooting Devices page (**Observability > Health > Troubleshoot Devices**). Based on the assessment of the test results, alerts and alarms are generated for the device and listed on the *Device-Name* page. If there are no alerts or alarms, the device status is healthy. The alerts and alarms are categorized as follows:

- Urgent Action Needed (Critical)
- Action Needed (Major)
- Being monitored (Minor)

See ["About the Events Page" on page 470](#) for a description of the alerts and alarms.

You can view a summary of the number of alerts and alarms for all the devices in an organization at the top of the Put Devices into Service page. The summary of the number of alerts and alarms makes it easy to find out which devices are functioning well, and which ones need attention.

The *Device-Name* page lists the test results and data collected from the respective device under the several accordions. You can expand the accordions to view details of the collected data by clicking the arrow present at the left of the accordion. You can also contact technical support by clicking the **Tech Support** link and view documentation by clicking the **View Documentation** link. The following accordions are present on the *Device-Name* page:

- Identity and Location—View general information about the device, the location where the device is installed and the trust score of the device. You can edit the site to which the device is assigned. See ["Identity and Location Data of a Device" on page 270](#).
- Remote Management—Displays the result of checks made on the management connection between the device and Routing Director. For information, See ["Remote Management Data and Test Results" on page 272](#).
- Hardware—View the temperature of the chassis and key performance indicators (KPIs) of all the hardware components and confirm that there are no alerts and alarms from any of the hardware components. Any alerts and alarms related to chassis hardware are also listed. For information on the test results and the hardware data displayed, See ["Hardware Data and Test Results" on page 278](#).
- Interfaces—View the power transmitted and received for optical pluggables and other data related to incoming and outgoing traffic at the interfaces. Any alerts and alarms related to the interfaces are also listed. See ["Interfaces Data and Test Results" on page 287](#).
- Software—View data such as vendor, software version, device model, SIRT advisories, and so on related to the software installed on the device. See ["Software Data and Test Results" on page 306](#).
- Configuration—View the configuration version, compliance of the committed configuration with Center for Internet Security (CIS) standards, and compliance score of the configuration. See ["Configuration Data and Test Results" on page 308](#).
- Routing and MPLS—View data related to the routing protocols. See ["Routing and MPLS Data and Test Results" on page 310](#).
- Connectivity—View data related to connectivity of the device with neighbors, Internet endpoints, cloud providers, and edge devices. See ["Device Connectivity Data and Tests Results" on page 343](#).

After all the checks are completed and if no issues are found, you can move the device to production and allow traffic to flow through the device.

To enable the device to carry production traffic, on top-left corner of the page, under the hostname, click **Put Into Service** and select **Put Into Service Now**. The Status of the device changes to In Service. The device is now deployed in the network and the device can allow the flow of live traffic.

You can continue monitoring the device from this page for details of any alerts and alarms that might be raised during production.

RELATED DOCUMENTATION

[View Network Topology Details | 504](#)

Identity and Location Data of a Device

Routing Director displays the identity and location information of a device besides the trust score of the device. [Table 64 on page 270](#) lists the data displayed in the Identity and Location accordion.

You can view the overall compliance of a device's hardware, software, and configuration with the rules in the trust plan, at the top-right corner of the accordion:

- **Healthy:** The device is compliant with the rules defined in the trust plan.
- **Being Monitored:** The device is being monitored for non-compliance with the rules in the trust plan.
- **Action Needed:** The device is low on trustworthiness and is vulnerable. You must intervene to ensure compliance.
- **Urgent Action Needed:** The device is very low on trustworthiness and is very vulnerable. You must intervene immediately to ensure compliance.

Table 64: Fields on the Identity and Location Accordion

Field	Description
Hostname	Hostname of the device.
Vendor	Vendor of the device.
Model	Model of the device; for example, ACX7100-48L.
Serial Number	Serial number of the device.

Table 64: Fields on the Identity and Location Accordion (*Continued*)

Field	Description
Management IP	Management IP address assigned to the device.
Score	Displays the most recent trust score of the device. Click the score to view all the trust scores for the device, measured at different intervals. An up arrow or down arrow next to the score indicates the rise or fall in the trust score as a percentage from the previous week's score. A lower score indicates that the device hardware, software, or configuration doesn't comply with one or more rules of the trust plan. NOTE: Compliance scanning is not supported for third-party devices. As a result, the compliance trust score for third-party devices is not available on the Identity and Location accordion.
Tags	Displays the tags assigned to the device. Click the Add Tag button to add tags to the device. In the Manage Tag page that appears, enter the tag in the <i>key:value</i> format. For example, <i>interfacetype:Internet</i>. A <i>key</i> represents a property or a parameter of the device or interface. A key must start with a letter and can be followed by lowercase letters, numbers, periods (.), hyphens (-), or underscores (_), and between 1 and 200 characters long. A <i>value</i> represents the value for that property or parameter. A value can include alphanumeric characters, periods (.), hyphens (-), and underscores (_) and between 1 and 200 characters long.
Site	Site where the device is installed. If you have superuser permissions, you will see the Edit (pencil) icon to edit the site. Click the Edit link to edit the site. See "Edit and Delete Sites" on page 91. Click the <i>site-name</i> to view all the devices present at the site on the Devices At Site <i>Site-Name</i> page. The Devices At Site <i>Site-Name</i> page is similar to the Put Devices into Service page, and you can perform all tasks (except moving the device to production) as from the Put Devices into Service page.

Table 64: Fields on the Identity and Location Accordion (*Continued*)

Field	Description
Location	Address of the site where the device is installed. If you have permissions, you can click the <i>address</i> to edit the address on the Edit Sites page.
Relevant Events	Lists the latest two alerts related to changes in the compliance score of the device in the order of severity. Hover over View Details to view details of that alert. Click View all Relevant Events to view all the alerts raised during the past seven days, related to the trust score on the Events for <i>Device-Name</i> page.

RELATED DOCUMENTATION

[Organization and Sites Overview | 57](#)

[Trust and Compliance Overview | 603](#)

Remote Management Data and Test Results

The Remote Management accordion provides details about the management connection between the device and Routing Director.

Routing Director displays the following information about the remote management accordion:

- The last time when the device successfully established or terminated an outbound SSH session with Routing Director.
- The last time when Routing Director received a system log message from the device.
- The last time when Routing Director received an alarm from the device.
- The last time when the device successfully established a gNMI session with Routing Director or when the session got terminated.
- The synchronization status between the device clock and the Network Time Protocol (NTP) server.

You can also release the device from the management of Routing Director from this accordion.



NOTE: For third-party devices, only gNMI-related information is available and supported. You can also stop Routing Director from managing the third-party devices by using the release device functionality.

For more information on the parameters displayed for all the devices in the accordion, see [Table 65 on page 273](#).



NOTE: All the fields display Data is not available when data is not collected for the remote management connection.

Table 65: Remote Management Accordion Data and Actions

Parameter	Description
Outbound SSH	Displays the date and time when the device successfully established an outbound SSH session with Routing Director or when the session got terminated. Hover over the timestamp to view the possible states. The states are: • Connected: The device has established an outbound SSH session with Routing Director. The connected state is also Indicated by a green circle with a white tick. • Disconnected: The outbound SSH session that the device established with Routing Director is terminated. The disconnected state is also indicated by a red circle with a white cross.

Table 65: Remote Management Accordion Data and Actions *(Continued)*

Parameter	Description
Syslog	Displays the date and time when the system log message from the device was last received by Routing Director. Hover over the timestamp to view details on the latest system log generated by the device. The details displayed are: - Severity—Severity level of the log message. The levels can be: - Critical - Error - Warning - Timestamp—Date and time when the device generated the system log message. - Appname—Application on the device that generated the log message. - Raw Message—Unprocessed system log message generated by the device. The unprocessed message contains additional log information such as, date and time when the message was generated, process and the ID of the process that generated the log message, and the device that generated the log message. - Org ID—Identifier of the organization to which the device belongs. - Host—Host name of the device. - Message—Processed message sent by the device, without any additional log information. Click the timestamp link to view additional details about all system logs generated in the organization from the Device Logs tab on the Events page (Observability > Health > Events > Device Logs).

Table 65: Remote Management Accordion Data and Actions *(Continued)*

Parameter	Description
Alarms	Displays the date and time when the alarm generated by the device was last received by Routing Director. Hover over the timestamp to view details about the latest alarm generated by the device. The details displayed are: • Device—Name of the device that generated the alarm. • Description—Description of the latest alarm raised on the device. • Last Received Time—Date and time when the latest alarm notification was received from the device. Click the timestamp link to view additional details about all alarms generated in the organization from the Alarms tab on the Events page (Observability > Health > Events > Alarms).
gNMI	Displays the date and time when the device successfully established or terminated a gNMI session with Routing Director. Hover over the timestamp to view the possible states. The states are: • Connected: The device has established a gNMI session with Routing Director. • Disconnected: The gNMI session that the device established with Routing Director got terminated.

Table 65: Remote Management Accordion Data and Actions *(Continued)*

Parameter	Description
Clock (NTP)	Displays whether the connection between the device and the NTP server is synchronized or not. The states are: • Synchronized: The device clock and the NTP server are synchronized. • Not Synchronized: The device clock and the NTP server are not synchronized. Click the link to view detailed history of synchronization between the NTP server and the device clock. The details displayed are: • Time—Date and time when the synchronization between the device and the NTP server was last tested. • Reference—IP address of the NTP server used as reference for synchronizing the clock on the device. If the IP address is unknown, then the field displays 0.0.0.0. • Status—Details about the synchronization including leap second measure, the current synchronization state, and so on. For more information about the NTP status, see Show NTP Status. • Time Offset—Difference in time between the NTP server and the device, before synchronization.

Table 65: Remote Management Accordion Data and Actions *(Continued)*

Parameter	Description
Release <i>Device</i>	By releasing a device, you stop Routing Director from managing the device. You can release a device when: • The device is no longer in use. • You want to reuse the device in another role or in another network. • You want to replace the device with another device. You can release a device by: • Clicking Release Device in this accordion or the Inventory page (Inventory > Devices > Network Inventory). You must be a user with the Super User role to release a device from this accordion. When you release a device by using this option, all the device configurations are retained on the device but the outbound SSH configuration on the device is deleted. Without the outbound SSH connection, the device is disconnected from Routing Director. You must manually delete the device configurations. • You can also release a device by removing the device from the Network Implementation Plan (Inventory> Device Onboarding > Network Implementation Plan). When you release a device by using the network implementation plan, all the device configurations that were committed on the device through the plan are deleted, but the outbound SSH connection is retained. You must manually delete the outbound SSH configuration. Alternatively, if you want to release all the devices that are part of the plan, you can delete the Network Implementation Plan (known as offboarding) that is used to onboard and manage the devices. NOTE: Deleting a Network Implementation Plan impacts all the devices that are part of the plan. For more information about releasing devices by using a Network Implementation Plan, see "Offboard a Network Implementation Plan" on page 259.

RELATED DOCUMENTATION

[Device Onboarding Workflow | 173](#)

[Device Management Workflow | 355](#)

Hardware Data and Test Results

SUMMARY

This topic provides information about the results of the tests that Routing Director executes to determine the health and functioning of the device hardware.

IN THIS SECTION

- [Overview | 278](#)
- [Hardware Details for *Device-Name* Page | 282](#)

Overview

The Hardware accordion displays the hardware data and results of the tests that Routing Director executes. These tests determine the health and functioning of a device's hardware. You can also view events (alerts and alarms), if any, for the device on the Hardware accordion and on the Hardware Details for *Device-Name* page.

To access the Hardware accordion, navigate to the **Observability > Health > Troubleshoot Devices** page. Click a device name to access the *Device-Name* page. Click the **Hardware (accordion)** in the Overview tab.

The top-right corner of the accordion displays the overall health of the device's hardware. The various states are:

- **Healthy**—The device's hardware (PSUs, fans, line cards, CPU, and memory) and temperature (of the Routing Engine, Routing Engine CPU, PSM, and chassis) is healthy.
- **Being Monitored**—The health of the device is being monitored.
- **Action Needed**—The device's hardware and temperature have issues that you must address.
- **Urgent Action Needed**—The device's hardware and temperature have issues that must be addressed immediately.

[Table 66 on page 279](#) lists the results of the hardware tests.



NOTE: Hardware tests are not supported on third-party devices. As a result, the hardware data and test results for third-party devices are not available on the Overview Tab.

Table 66: Results of Hardware Tests

Field	Description
PSUs	Total number of power supply units (PSUs) present in the device and the total number of unhealthy PSUs. A PSU is marked unhealthy when: • The supply exceeds the high and low threshold limits. • The PSU temperature exceeds the high and low threshold limits. Click the link next to PSUs to view the threshold limits and the performance of the PSUs for a week, a day, 3 hours, 1 hour, 30 minutes, or a custom time period. See "Hardware Details for Device-Name Page" on page 282 for more information.
Fans	Total number of fans present in the device and the total number of unhealthy fans. A fan is marked unhealthy when the RPM exceeds the high and low threshold limits. Click the link next to Fans to view the threshold limits and the performance of the fans for a week, a day, 3 hours, 1 hour, 30 minutes, or a custom time period. See "Hardware Details for Device-Name Page" on page 282 for more information.
Linecards	Total number of line cards in the device and the total number of unhealthy line cards. A line card is marked unhealthy when the KPIs defined for that line card is not met. Click the link next to Linecards to view the threshold limits and the performance of the line cards for a week, a day, 3 hours, 1 hour, 30 minutes, or a custom time period. See "Hardware Details for Device-Name Page" on page 282 for more information. NOTE: Line card charts are not available on some ACX Series, PTX Series, and MX Series devices as the flexible PIC concentrator (FPC) fields are not supported on these devices. See Table 68 on page 285 for more information.

Table 66: Results of Hardware Tests *(Continued)*

Field	Description
CPU	Total number of CPUs in the device and the total number of unhealthy CPUs. A CPU is marked unhealthy when the CPU utilization exceeds the threshold limit. Click the link next to CPU to view the threshold limits and the performance of the CPU for a week, a day, 3 hours, 1 hour, 30 minutes, or a custom time period. See "Hardware Details for Device-Name Page" on page 282 for more information.
Memory	Memory utilized by Routing Engines and line cards, and the total number of unhealthy memory units. Device memory is marked unhealthy when the memory runs low or is insufficient. Click the link next to Memory to view the threshold limits and memory utilization of Routing Engines for a week, a day, 3 hours, 1 hour, 30 minutes, or a custom time period. See "Hardware Details for Device-Name Page" on page 282 for more information.
Temperature	Routing Engine temperature, PSM temperature, Routing Engine CPU temperature, line card temperature, and chassis temperature in degree Celsius. Temperature is marked unhealthy when the temperature exceeds the high and low threshold limits. Click the link next to Temperature to view more information on temperature utilization, which is displayed over a period of a week, a day, 3 hours, 1 hour, 30 minutes, or a custom time period. See "Hardware Details for Device-Name Page" on page 282 for more information.
Authenticity	Authenticity of the device hardware. Genuine Juniper Hardware is displayed if the device is a Juniper device.
End of Support	End of support information of the device.

Table 66: Results of Hardware Tests *(Continued)*

Field	Description
SIRT Advisories	Total number of Security Incident Resource Team (SIRT) advisories for the device and the software running on the device. Click the link next to SIRT Advisories to view the list of vulnerabilities that affect the device, and the software installed on the device, which is displayed on the Trust > Vulnerabilities page.
Relevant Events	Displays two issues or anomalies related to the hardware in order of severity. Hover over View Details to view more information about an issue in a pop-up. Click View All Relevant Events to view all hardware issues present on the device, on the Events for <i>Device-Name</i> page. You can view relevant events for the past seven days.
Show LEDs, Ports & Cables on Chassis	Click the Show LEDs, Ports, Cables on Chassis toggle button to view or hide the device chassis. Hover over the CPU, memory, fans, power, and temperature icons to view a snapshot of the performance of each component. Click the Port Status drop-down list to view: • Show All (default option) • Show Up • Show Down • Show None You can zoom in, zoom out, and reset a device chassis.



NOTE: KPI-related data may not be displayed on the accordions and on the graphs when:

- There are device errors
- The data displayed is derived from or is dependent on another KPI that is not being streamed

Hardware Details for *Device-Name* Page

To access the Hardware Details on the Routing Director GUI, click **Observability > Health > Troubleshoot Devices > *Device-Name* > Overview > Hardware (accordion)**. Click the *link* next to each hardware component to view its details.

You can view the health and performance of the device hardware components on the Hardware Details for *Device-Name* page.

The nine accordions on this page provide information on the health and functioning of the hardware components and temperature. [Table 67 on page 282](#) describes the accordions.

Table 67: Accordions on the Hardware Details for *Device-Name* Page

Accordion	Description
PSUs	Select PSM Power or PSM temperature from the Show PSUs drop-down list to view a list of up to six PSUs. These PSUs are listed in order of severity of the events that have occurred on them. The PSU, with the most critical events, appear at the top of the list. NOTE: Power supply module (PSM) temperature information is not available for MX204, MX480, and MX960 devices. Click the toggle button next to a PSU in the Show PSUs list to view the performance of the PSU in a graph. Alerts related to PSM Power or PSM Temperature are displayed above the graph. See "Performance Graphs" on page 285 for more information.
PSUs State	View the states of all PSUs on a graph. Alerts, if any, related to the PSUs are also displayed. See "Performance Graphs" on page 285 for more information.
Fans	View a list of up to six fans and information related to the speeds of the fan, in the order of severity of events that have occurred on them. The fan with the most critical events appears at the top of the list. NOTE: Charts related to the speed of the fan (rpm-percent) is not available for MX480, MX960, MX10004, and MX10008 devices. Click the toggle button next to a fan in the Show FAN Speeds list to view the performance of the fan in a graph. Fan state (alerts related to rpm-percent) is displayed above the graph. See "Performance Graphs" on page 285 for more information.

Table 67: Accordions on the Hardware Details for *Device-Name* Page (Continued)

Accordion	Description
Fans State	View the states of all fans on a graph. Alerts, if any, related to the fans are also displayed. See "Performance Graphs" on page 285 for more information.
Linecard	Select any option (Temperature line cards, Line cards CPU, Line cards Memory) from the Show Linecards drop-down list to view a list of up to six line cards. These line cards are listed in the order of severity of the events that have occurred on them. The line card with the most critical events appears at the top of the list. Click the toggle button next to a line card in the Show Linecards list to view the performance and alerts of the line card in a graph. See "Performance Graphs" on page 285 for more information. NOTE: - Alerts related to pfe-data-error-discard, pfe-bad-route-discard, pfe-bits-to-test-discard, pfe-fabric-discards, pfe-info-cell-discard, pfe-invalid-iif-discard, pfe-nexthop-discard, pfe-stack-overflow-discard, pfe-stack-underflow-discard, and pfe-tcp-header-error-discard are displayed. - Line card charts are not available on some ACX Series, PTX Series, and MX Series devices as the flexible PIC concentrator (FPC) fields are not supported on these devices. See Table 68 on page 285 for more information.
Linecards State	View the states of all line cards on a graph. Alerts, if any, related to the line cards are also displayed. See "Performance Graphs" on page 285 for more information.

Table 67: Accordions on the Hardware Details for *Device-Name* Page (Continued)

Accordion	Description
CPU	Select Routing Engines or Line cards CPU from the Show CPU Utilization drop-down list to view CPU utilization of up to six Routing Engines or line cards. These Routing Engines and line cards are listed in the order of severity of the events that have occurred on them. The CPU with the most critical events appears at the top of the list. NOTE: CPU utilization information is not available for PTX10002 device. Click the toggle button next to a CPU in the Show CPU Utilization list to view the utilization and alerts of the CPU in a graph. See "Performance Graphs" on page 285 for more information.
Memory	Select Routing Engines or Line cards Memory from the Show Memory Utilization drop-down list to view memory utilization of up to six Routing Engines or line cards. These Routing Engines and line cards are listed in the order of severity of the events that have occurred on them. The memory unit with the most critical events appears at the top of the list. Click the toggle button next to a memory unit in the Show Memory Utilization list to view the performance and alerts of the memory unit in a graph. See "Performance Graphs" on page 285 for more information.
Temperature	Select Routing Engines, Routing Engines CPU, Chassis, Temperature Line cards or PSM Temperature from the Show Temperature drop-down list to view temperature of up to six hardware components. These components are listed in the order of severity of the events that have occurred on them. The component with the highest temperature is listed at the top of the list. Device chassis temperature is displayed in degree Celsius. NOTE: Chassis temperature information is not available for MX204, MX240, MX304, MX10004, and MX10008 devices. Click the toggle button next to a component in the Show Temperatures list to view the temperature utilization of that component in a graph. See "Performance Graphs" on page 285 for more information.

Table 68: Line Card Charts Support

Device Family	Device Series	FPC Fields Not Supported
ACX Series	ACX7100-32C, ACX7100-48L, ACX7024, ACX7024X, ACX7509, ACX7348	fpc-temperature, fpc-cpu-utilization, fpc-buffer-memory-utilization
MX Series	MX204, MX240, MX304, MX480, MX960, MX10004, MX10008	fpc-temperature, fpc-cpu-utilization
PTX Series	PTX10002	fpc-temperature, fpc-cpu-utilization, fpc-buffer-memory-utilization

Performance Graphs

The graphs on the Hardware Details for *Device-Name* page display the performance of the hardware components. You can also view information on alerts and breaches, if any, on these graphs. [Figure 20 on page 285](#) shows the graphs for fans in a device.

Figure 20: Fans Accordion

The fans present in the device are listed on the left of the Fan accordion, in the order of severity of events that have occurred. You can view up to six fans at a time with the fan that is in the most critical state displayed at the top of the list. To view fans that are not listed, click the **Select Fans** drop-down list and select the fans. However, you must clear a previously selected fan to be able to select another fan.

Click the toggle button next to the name of the fan, to view the performance of the fan in a graph. The graph displays two lines showing the high (in orange) and critical (in red) threshold levels related to the speed of the fan. You can also zoom into a particular portion of the graph to view more information about events that have occurred.

To view the performance of more than one fan on the graph, click the toggle button next to the name of the fan in the Show Fan Speeds list. Details of the fans displayed on the graph are also listed in a table below the graph as shown in [Figure 20 on page 285](#). You can also click the option buttons on the left of a fan name in the table to highlight the graph for that fan.



NOTE: The toggle button next to the name of the component and the drop-down list is displayed only when there are more than six components available.

You can view the performance of a fan for a week, a day, 3 hours, 1 hour, 30 minutes, or a custom time period. By default, performance for the past 30 minutes is displayed. To change this period, click the **Week, Day, 3 hrs, 1 hr, 30 mins, or Custom** buttons provided above the graph.

You can view more than 25 data points on a graph related to events (in real time) that have occurred on the fan when you select the 30-minute time period. However, you can only view up to 25 data points related to events when you select a week, a day, 3 hours, 1 hour, or a custom time period (of more than 30 minutes). Data is aggregated to ensure that not more than 25 data points are plotted on the graph at once.

The graph auto-refreshes at an interval depending on the time range for which the information is displayed. See [Table 69 on page 287](#) for more information. However, you can refresh the graph at any point by clicking the **Refresh** icon provided above the graph.

You can also click the pop-out button next to the graph to open the graph in a new tab. You can view all customizations that you made on the graph in the parent tab, in the new tab.

The most critical alert (issues and anomalies) is displayed just above the graph and next to the quick chart. To view other alerts, click the [link](#) just below the alert. The quick chart displays the performance of the fan that you selected from the **Select Fans** drop-down list. However, if alerts related to the performance of any fan is generated, the fan with the most severe alert is displayed on the quick chart by default.

Alerts, if any, related to the fan is also displayed on the graph, and in the table below the graph. You can also open the graph in a new tab. When you open a graph in a new tab, you can view the following information in the new tab as well:

- Alerts related to the fan.

Alerts are refreshed across all open tabs simultaneously, when:

- An alert is refreshed in any one of the open tabs.

- The last alert fetched was beyond three minutes.
- List of fans that you toggled to view from the Select Fans drop-down list.
- Fan that you selected from the table below the graph.

Table 69: Auto-Refresh Rate

Time Range	Auto-Refreshed
Weekly	Every 16.8 hours
Daily	Every 58 minutes
Every three hours	Every 8 minutes
Hourly	Every 3 minutes
Thirty minutes	Every 2 minutes
Custom	No auto-refresh

You can similarly view the graphs and alerts related to the performance of other hardware components and temperature.

Interfaces Data and Test Results

SUMMARY

This topic provides information about the results of the tests that Routing Director executes to determine the state of the device interfaces.

IN THIS SECTION

- [Overview | 288](#)
- [Pluggables Details for *Device-Name* Page | 292](#)
- [Input Traffic Details for *Device-Name* Page | 295](#)
- [Output Traffic Details for *Device-Name* Page | 299](#)

Overview

The Interfaces accordion displays the interfaces data and results of the tests that Routing Director executes to determine that:

- The device's interfaces are up.
- There are no port flapping issues.
- The input and output traffic does not exceed the threshold limit.

You can also view events (alerts and alarms), if any, for the device on the Interfaces accordion.

To access the Interfaces accordion, navigate to the **Observability > Health > Troubleshoot Devices** page. Click a device name to access the *Device-Name* page. Click the **Interfaces (accordion)** in the Overview tab.

The top-right corner of the accordion displays the overall health of the interfaces. The various states are:

- **Healthy**—The interfaces are healthy.
- **Being Monitored**—The health of the interfaces is being monitored.
- **Action Needed**—The interfaces have issues that you must address (might not be immediately).
- **Urgent Action Needed**—The interfaces have issues that you must address immediately.

[Table 70 on page 289](#) lists the results of the interface checks.



NOTE: Interfaces tests are not supported on third-party devices. As a result, interfaces data and test results for third-party devices are not available on the Overview Tab.

Table 70: Results of Interface Checks

Field	Description
Pluggables	Total number of available pluggables, total number of monitored pluggables, and the total number and names of the unhealthy pluggables. • Total number of available pluggables—Total number of pluggables on the device. • Total number of monitored pluggables—Total number of pluggables on the device that are monitored. • Total number of unhealthy pluggables—Total number of pluggables that are monitored and are unhealthy. A pluggable is marked unhealthy when the optical temperature and the optical signal strength (Tx power, Rx power) KPIs exceed the threshold limits. Click the link next to Pluggables to view: • Details about optical temperature • Power of the signal leaving the device • Power of the incoming signal received from the neighboring device You can view this information for the past week, day, 3 hours, 1 hour, 30 minutes, or a custom time period. See "Pluggables Details for Device-Name Page" on page 292 for more information.

Table 70: Results of Interface Checks *(Continued)*

Field	Description
Input Traffic	Total number of available interfaces, total number of monitored interfaces, and the total number and names of the unhealthy interfaces. • Total number of available interfaces—Total number of interfaces on the device. • Total number of monitored interfaces—Total number of interfaces on the device that are monitored. • Total number of unhealthy interfaces—Total number of interfaces that are monitored and are unhealthy. Unhealthy interfaces are those interfaces with errors such as CRC errors, FEC uncorrected errors, and other input errors. Click the link next to Input Traffic to view: • Details about signal functionality (Rx signal loss) • Optical Rx power You can view this information for the past week, day, 3 hours, 1 hour, 30 minutes, or a custom time period. See "Input Traffic Details for Device-Name Page" on page 295 for more information.
Output Traffic	Total number of available interfaces, total number of monitored interfaces, and the total number and names of the unhealthy interfaces. • Total number of available interfaces—Total number of interfaces on the device. • Total number of monitored interfaces—Total number of interfaces on the device that are monitored. • Total number of unhealthy interfaces—Total number of interfaces that are monitored and are unhealthy. Unhealthy interfaces are those interfaces with errors such as CRC errors, FEC uncorrected errors, and other output errors. Click link next to Output Traffic to view: • Details about signal functionality (Tx signal loss, and Tx laser disabled alarm) • Optical power of the outgoing signal You can view this information for the past week, day, 3 hours, 1 hour, 30 minutes, or a custom time period. See "Output Traffic Details for Device-Name Page" on page 299 for more information.

Table 70: Results of Interface Checks *(Continued)*

Field	Description
Interfaces	Total number of available interfaces, total number of monitored interfaces, the total number of interfaces that are down, and the total number of interfaces with link flapping issues. • Total number of available interfaces—Total number of interfaces on the device. • Total number of monitored interfaces—Total number of interfaces on the device that are monitored. • Total number of unhealthy interfaces—Total number of interfaces that are monitored and are unhealthy. The names of the interfaces that are down and with link flapping issues are also displayed. An interface is marked unhealthy when the interface: • Is operational and has errors. • Is not operational. • Traffic exceeds the high and low threshold limits. • Has link flapping issues. Click link next to Interfaces to view details about link state and port flapping issues. You can view this information for the past week, day, 3 hours, 1 hour, 30 minutes, or a custom time period. See "Interfaces Details for Device-Name Page" on page 303 for more information.
Relevant Events	Displays two issues or anomalies with respect to the interfaces in order of severity. Hover over View Details to view more information about an issue in a pop-up. Click View All Relevant Events to view all device interface issues, on the Events for <i>Device-Name</i> page. You can view relevant events for the past seven days.



NOTE: KPI-related data may not be displayed on the accordions and on the graphs when:

- There are device errors
- The data displayed is derived from or is dependent on another KPI that is not being streamed

Pluggables Details for *Device-Name* Page

To access the Pluggables Details on the Routing Director GUI, click **Observability > Health > Troubleshoot Devices > *Device-Name* > Overview > Interfaces (accordion) > Pluggables**. Click the *link* next to each components to view its details.

You can view the health and functioning of the pluggables from the Pluggables Details for *Device-Name* page.

The three accordions on this page provide information on optical temperature, transmission power, and receiving power. [Table 71 on page 292](#) describes the accordions.

Table 71: Accordions on the Pluggables Details for *Device-Name* Page

Accordion	Description
Optical Temperature	View the optical temperature for the optical interfaces. You can view up to six interfaces on a graph at a time. These interfaces are listed in the order of severity of the events that have occurred on them. The interface with the most critical events appears at the top of the list. Click the toggle button next to an interface in the Show Interfaces list to view the performance and alerts of the interface in a graph. See "Performance Graphs" on page 293 for more information.
Optical Tx Power	View the outgoing signal strength for the optical interfaces. You can view up to six interfaces on a graph at a time. These interfaces are listed in the order of severity of the events that have occurred on them. The interface with the most critical events appears at the top of the list. Click the toggle button next to an interface in the Show Interfaces list to view the performance and alerts of the interface in a graph. See "Performance Graphs" on page 293 for more information.


Table 71: Accordions on the Pluggables Details for *Device-Name* Page (Continued)

Accordion	Description
Optical Rx Power	View the incoming signal strength for the optical interface. You can view up to six interfaces on a graph at a time. These interfaces are listed in the order of severity of the events that have occurred on them. The interface with the most critical events appears at the top of the list. Click the toggle button next to an interface in the Show Interfaces list to view the performance and alerts of the interface in a graph. See "Performance Graphs" on page 293 for more information.

Performance Graphs

The graphs on the Pluggables Details for *Device-Name* page display the health and functioning of the pluggables. You can also view information on alerts and breaches, if any, on these graphs. [Figure 21 on page 293](#) shows details about optical temperature for a device interface on the graph.

Figure 21: Optical Temperature Accordion



The interfaces present in the device are listed on the left of the Optical Temperature accordion, in order of severity of events that have occurred on them. You can view up to six interfaces at a time. To view interfaces that are not listed, click the **Select Interfaces** drop-down list and select the interface. However, you must clear a previously selected interface to be able to select another interface.

Click the toggle button next to the name of the interface to view the performance of that interface in a graph. The graph displays two lines showing the high (in orange) and critical (in red) threshold levels

related to the events that have occurred on the interface. Another line on the graph displays the data points related to events that have occurred on the interface. The color of the data points depends on the color of the interface that you have selected. The color of the interface is assigned automatically. Click any data point to view more information about the corresponding event in a pop-up. The cause for the alert, if any, is also displayed.

To view the performance of more than one interface on the graph, click the toggle button next to the name of the interface in the Show Interfaces list. Details of the interfaces displayed on the graph are also listed in a table below the graph as shown in [Figure 21 on page 293](#). You can also click the option buttons on the left of an interface name in the table to highlight the graph for that interface.

You can view the optical temperature information for a week, a day, 3 hours, 1 hour, 30 minutes, or a custom time period. By default, information for the past 30 minutes is displayed. To change this period, click the **Week**, **Day**, **3 hrs**, **1 hr**, **30 mins**, or **Custom** buttons provided above the graph.

The graph auto-refreshes at an interval depending on the time range for which the information is displayed. See [Table 72 on page 295](#) for more information. However, you can refresh the graph at any point by clicking the **-Refresh** icon provided above the graph.

You can also click the pop-out button next to the graph to open the graph in a new tab and view all customizations that you made on the graph in the parent tab, in the new tab.

The most critical alert (issues and anomalies) related to optical temperature is displayed just above the graph and next to the quick chart. To view other alerts, click the [*link*](#) just below the alert. The quick chart displays the performance of the interface that you selected from the **Select Interfaces** drop-down list. However, if alerts related to the performance of any interface is generated, the interface with the most severe alert is displayed on the quick chart by default.

Alerts, if any, related to optical temperature is displayed on the graph, and in the table below the graph. You can also open the graph in a new tab. When you open a graph in a new tab, you can view the following information in the new tab as well:

- Alerts related to optical temperature.

Alerts are refreshed across all open tabs simultaneously, when:

- An alert is refreshed in any one of the open tabs.
- The last alert fetched was beyond three minutes.
- List of interfaces that you toggled to view from the Select Interfaces drop-down list.
- Interface that you selected from the table below the graph.

Table 72: Auto-Refresh Rate

Time Range	Auto-Refreshed
Weekly	Every 16.8 hours
Daily	Every 58 minutes
Every three hours	Every 8 minutes
Hourly	Every 3 minutes
Thirty minutes	Every 2 minutes
Custom	No auto-refresh

You can similarly view the graphs and alerts related to strength of the outgoing signal (Optical Tx Power), and strength of the incoming signal (Optical Rx Power).

Input Traffic Details for *Device-Name* Page

To access the Input Traffic Details on the Routing Director GUI, click **Observability > Health > Troubleshoot Devices > *Device-Name* > Overview > Interfaces (accordion) > Input Traffic**. Click the *link* next to each traffic component to view its details.

You can view information about input traffic flow on the Input Traffic Details for *Device-Name* page.

The eight accordions on this page provide information about signal functionality, the highest and lowest power of the incoming signal, receiving (Rx) power, input traffic range, input errors, CRC errors, and FEC Uncorrected Errors.

[Table 73 on page 296](#) describes the accordions.

Table 73: Accordions on the Input Traffic Details for *Device-Name* Page

Accordion	Description
Signal Functionality	View signal functionality (receiving [Rx] signal loss) at the device's interfaces. You can view data for up to six interfaces on a graph at a time. These interfaces are listed in the order of severity of the events that have occurred on them. The interface with the most critical events appears at the top of the list. Click the toggle button next to an interface in the Rx Loss Of Signal Alarm list to view the performance and alerts of the interface in a graph. See "Performance Graphs" on page 297 for more information.
Optical Rx Power	View optical power of the incoming signal (Rx power) at the device's interfaces. You can view data for up to six interfaces on a graph at a time. These interfaces are listed in the order of severity of the events that have occurred on them. The interface with the most critical events appears at the top of the list. Click the toggle button next to an interface in the Show Interfaces list to view the performance and alerts of the interface in a graph. See "Performance Graphs" on page 297 for more information.
Input Traffic	View input traffic at the device's interfaces. You can view data for up to six interfaces on a graph at a time. These interfaces are listed in the order of severity of the events that have occurred on them. The interface with the most critical events appears at the top of the list. Click the toggle button next to an interface in the Show Input Traffic list to view the performance and alerts of the interface in a graph. See "Performance Graphs" on page 297 for more information.
Input Errors	View input errors generated at the device's interfaces. You can view data for up to six interfaces on a graph at a time. These interfaces are listed in the order of severity of the events that have occurred on them. The interface with the most critical events appears at the top of the list. Click the toggle button next to an interface in the Show Input Errors list to view the performance and alerts of the interface in a graph. See "Performance Graphs" on page 297 for more information.

Table 73: Accordions on the Input Traffic Details for *Device-Name* Page (Continued)

Accordion	Description
FEC Uncorrected Errors	View forward error correction (FEC) uncorrected errors generated at the device's interfaces. You can view data for up to six interfaces on a graph at a time. These interfaces are listed in the order of severity of the events that have occurred on them. The interface with the most critical events appears at the top of the list. NOTE: FEC uncorrected errors are available on interfaces that support speeds equal to or greater than 100-Gbps. Click the toggle button next to an interface in the Show Interfaces list to view the performance and alerts of the interface in a graph. See "Performance Graphs" on page 297 for more information.
CRC Errors	View cyclic redundancy check (CRC) errors generated at the device's interfaces. You can view data for up to six interfaces on a graph at a time. These interfaces are listed in the order of severity of the events that have occurred on them. The interface with the most critical events appears at the top of the list. Click the toggle button next to an interface in the Show CRC Errors list to view the performance and alerts of the interface in a graph. See "Performance Graphs" on page 297 for more information.
Framing Errors	View framing errors generated at the device's interfaces. You can view data for up to six interfaces on a graph at a time. These interfaces are listed in the order of severity of the events that have occurred on them. The interface with the most critical events appears at the top of the list. Click the toggle button next to an interface in the Show Interfaces list to view the performance and alerts of the interface in a graph. See "Performance Graphs" on page 297 for more information.


Performance Graphs

The graphs on the Input Traffic Details for *Device-Name* page display detailed information about input traffic flow. You can also view information about alerts and breaches, if any, on these graphs.

To view information related to input traffic on a graph, navigate to **Observability > Health > Troubleshoot Devices > *Device-Name* > Overview > Interfaces (accordion) > Input Traffic *data-link***, and click any accordion. Detailed information and graphs related to that accordion are displayed within that input traffic accordion.

[Figure 22 on page 298](#) shows details about input traffic flow for a device interface on the graph.

Figure 22: Input Traffic Accordion



The interfaces present in the device are listed on the left of the Input Traffic accordion, in order of severity of events that have occurred on them. You can view up to six interfaces at a time. To view interfaces that are not listed, click the **Select Interfaces** drop-down list and select the interface. However, you must clear a previously selected interface to be able to select another interface.

Click the toggle button next to the name of the interface, to view details on input traffic flow of that interface in a graph. The graph displays two lines showing the high (in orange) and critical (in red) threshold levels related to the events that have occurred on the interface. Another line on the graph displays the data points related to events that have occurred on the interface. The color of the data points depends on the color of the interface that you have selected. The color of the interface is assigned automatically. Click any data point to view more information about the corresponding event in a pop-up. The cause for the alert, if any, is also displayed.

To view the performance of more than one interface on the graph, click the toggle button next to the name of the interface in the Show Input Traffic list. Details of the interfaces displayed on the graph are also listed in a table below the graph as shown in [Figure 22 on page 298](#). You can also click the option buttons on the left of an interface name in the table to highlight the graph for that interface.

You can view information related to input traffic for a week, a day, 3 hours, 1 hour, 30 minutes, or a custom time period. By default, information for the past 30 minutes is displayed. To change this period, click the **Week**, **Day**, **3 hrs**, **1 hr**, **30 mins**, or **Custom** buttons provided above the graph.

The graph auto-refreshes at an interval depending on the time range for which the information is displayed. See [Table 74 on page 299](#) for more information. However, you can also choose to refresh the graph at any point by clicking the **Refresh** icon provided above the graph.

You can also click the pop-out button next to the graph to open the graph in a new tab and view all customizations that you made on the graph in the parent tab, in the new tab.

The most critical alert (issues and anomalies) related to input traffic is displayed just above the graph and next to the quick chart. To view other alerts, click the [link](#) just below the alert. The quick chart displays the performance of the interface that you selected from the **Select Interfaces** drop-down list. However, if alerts related to the performance of any interface is generated, the interface with the most severe alert is displayed on the quick chart by default.

Alerts, if any, on events related to input traffic that have occurred are displayed on the graph, and in the table below the graph. You can also open the graph in a new tab. When you open a graph in a new tab, you can view the following information in the new tab as well:

- Alerts related to input traffic.

Alerts are refreshed across all open tabs simultaneously, when:

- An alert is refreshed in any one of the open tabs.
- The last alert fetched was beyond three minutes.
- List of interfaces that you toggled to view from the Select Interfaces drop-down list.
- Interface that you selected from the table below the graph.

Table 74: Auto-Refresh Rate

Time Range	Auto-Refreshed
Weekly	Every 16.8 hours
Daily	Every 58 minutes
Every three hours	Every 8 minutes
Hourly	Every 3 minutes
Thirty minutes	Every 2 minutes
Custom	No auto-refresh

You can similarly view the graphs and alerts related to signal functionality, the highest and lowest power of the incoming signal, receiving (Rx) power, input errors, and CRC errors.

Output Traffic Details for *Device-Name* Page

To access the Output Traffic Details on the Routing Director GUI, click **Observability > Health > Troubleshoot Devices > *Device-Name* > Overview > Interfaces (accordion) > Output Traffic**. Click the [link](#) next to each hardware component to view its details.

You can view detailed information about output traffic flow from the Output Traffic Details for *Device-Name* page.

The seven accordions on this page provide information about signal functionality, the highest and lowest power of the outgoing signal, transmission (Tx) power, output traffic range, output errors, and CRC errors.

[Table 75 on page 300](#) describes the accordions.

Table 75: Accordions on the Output Traffic Details for *Device-Name* Page

Accordion	Description
Signal Functionality	View signal functionality (transmission [Tx] signal loss, and Tx laser disabled alarm) at the device's interfaces. You can view data for up to six interfaces on a graph at a time. These interfaces are listed in the order of severity of the events that have occurred on them. The interface with the most critical events appears at the top of the list. Click the toggle button next to an interface in the Show Signal Functionality list to view the performance and alerts of the interface in a graph. See "Performance Graphs" on page 301 for more information.
Optical Tx Power	View the power of the outgoing signal (Tx power) at the device's interfaces. You can view data for up to six interfaces on a graph at a time. These interfaces are listed in the order of severity of the events that have occurred on them. The interface with the most critical events appears at the top of the list. Click the toggle button next to an interface in the Show Interfaces list to view the performance and alerts of the interface in a graph. See "Performance Graphs" on page 301 for more information.
Output Traffic	View the output traffic at the device's interfaces. You can view data for up to six interfaces on a graph at a time. These interfaces are listed in the order of severity of the events that have occurred on them. The interface with the most critical events appear at the top of the list. Click the toggle button next to an interface in the Show Interfaces list to view the performance and alerts of the interface in a graph. See "Performance Graphs" on page 301 for more information.

Table 75: Accordions on the Output Traffic Details for *Device-Name* Page (Continued)

Accordion	Description
Output Errors	View the output errors generated at the device's interfaces. You can view data for up to six interfaces on a graph at a time. These interfaces are listed in the order of severity of the events that have occurred on them. The interface with the most critical events appears at the top of the list. Click the toggle button next to an interface in the Show Interfaces list to view the performance and alerts of the interface in a graph. See "Performance Graphs" on page 301 for more information.
FEC Uncorrected Errors	View the forward error correction (FEC) uncorrected errors generated at the device's interfaces. You can view data for up to six interfaces on a graph at a time. These interfaces are listed in the order of severity of the events that have occurred on them. The interface with the most critical events appears at the top of the list. NOTE: FEC uncorrected errors are available on interfaces that support speeds equal to or greater than 100-Gbps. Click the toggle button next to an interface in the Show Interfaces list to view the performance and alerts of the interface in a graph. See "Performance Graphs" on page 301 for more information.
Output CRC Errors	View the cyclic redundancy check (CRC) errors at the device's interfaces. You can view data for up to six interfaces on a graph at a time. These interfaces are listed in the order of severity of the events that have occurred on them. The interface with the most critical events appears at the top of the list. Click the toggle button next to an interface in the Show Interfaces list to view the performance and alerts of the interface in a graph. See "Performance Graphs" on page 301 for more information.

Performance Graphs

The graphs on the Output Traffic Details for *Device-Name* page display detailed information about output traffic flow. You can also view information about alerts and breaches, if any, on these graphs.

[Figure 23 on page 302](#) shows details about output traffic flow for a device interface on the graph.

Figure 23: Output Traffic Accordion



The interfaces present in the device are listed on the left of the Output Traffic accordion, in order of severity of events that have occurred on them. You can view up to six interfaces at a time. To view interfaces that are not listed, click the **Select Interfaces** drop-down list and select the interface. However, you must clear a previously selected interface to be able to select another interface.

Click the toggle button next to the name of the interface, to view details on output traffic flow of that interface in a graph. The graph displays two lines showing the high (in orange) and critical (in red) threshold levels related to the events that have occurred on the interface. Another line on the graph displays the data points related to events that have occurred on the interface. The color of the data points depends on the color of the interface that you have selected. The color of the interface is assigned automatically. Click any data point to view more information about the corresponding event in a pop-up. The cause for the alert, if any, is also displayed.

To view the performance of more than one interface on the graph, click the toggle button next to the name of the interface in the Show Interfaces list. Details of the interfaces displayed on the graph are also listed in a table below the graph as shown in [Figure 23 on page 302](#). You can also click the option buttons on the left of an interface name in the table to highlight the graph for that interface.

You can view information related to output traffic for a week, a day, 3 hours, 1 hour, 30 minutes, or a custom time period. By default, information for the past 30 minutes is displayed. To change this period, click the **Week**, **Day**, **3 hrs**, **1 hr**, **30 mins**, or **Custom** buttons provided above the graph.

The graph auto-refreshes at an interval depending on the time range for which the information is displayed. See [Table 76 on page 303](#) for more information. However, you can refresh the graph at any point by clicking the **Refresh** icon provided above the graph.

You can also click the pop-out button next to the graph to open the graph in a new tab and view all customizations that you made on the graph in the parent tab, in the new tab.

The most critical alert (issues and anomalies) related to output traffic is displayed just above the graph and next to the quick chart. To view other alerts, click the *link* just below the alert. The quick chart displays the performance of the interface that you selected from the **Select Interfaces** drop-down list. However, if alerts related to the performance of any interface are generated, the interface with the most severe alert is displayed on the quick chart by default.

Alerts, if any, on events related to output traffic that have occurred are displayed on the graph, and in the table below the graph. You can also open the graph in a new tab. When you open a graph in a new tab, you can view the following information in the new tab as well:

- Alerts related to output traffic.

Alerts are refreshed across all open tabs simultaneously, when:

- An alert is refreshed in any one of the open tabs.
- The last alert fetched was beyond three minutes.
- List of interfaces that you toggled to view from the Select Interfaces drop-down list.
- Interface that you selected from the table below the graph.

Table 76: Auto-Refresh Rate

Time Range	Auto-Refreshed
Weekly	Every 16.8 hours
Daily	Every 58 minutes
Every three hours	Every 8 minutes
Hourly	Every 3 minutes
Thirty minutes	Every 2 minutes
Custom	No auto-refresh

You can similarly view the graphs and alerts related to signal functionality, the highest and lowest power of the outgoing signal, transmission (Tx) power, output errors, and output CRC errors.

Interfaces Details for *Device-Name* Page

To access the Interfaces Details on the Routing Director GUI, click **Observability > Health > Troubleshoot Devices > Device-Name > Overview > Interfaces (accordion) > Interfaces**. Click the *link* next to each interface component to view its details.

You can view detailed information about link state performance and issues, and port flapping issues related to physical interfaces from the Interfaces Details for *Device-Name* page.

The two accordions on this page provide information about link state and port flapping issues related to physical interfaces.

[Table 77 on page 304](#) describes the accordions.

Table 77: Accordions on the Interfaces Details for *Device-Name* Page

Accordion	Description
Link State	View link state issues present at the device's interfaces. You can view up to six interfaces on a graph at a time. The interface with the most critical events appears at the top of the list. Click the toggle button next to an interface in the Show Link State list to view the performance and alerts of the interface in a graph. See "Performance Graphs" on page 304 for more information.
Link Flap	View information related to link flapping issues at the device's interfaces. You can view up to six interfaces on a graph at a time. The interface with the most critical events appears at the top of the list. Click the toggle button next to an interface in the Show Link Flaps list to view the performance and alerts of the interface in a graph. See "Performance Graphs" on page 304 for more information.

Performance Graphs

The graphs on Interfaces Details for *Device-Name* page display information about link state and port flapping issues related to physical interfaces. You can also view information about alerts and breaches, if any, on these graphs.

[Figure 24 on page 305](#) shows details about link state performance and issues for an interface on the graph.

Figure 24: Link State Accordion



The interfaces for the Link State accordion are listed on the left of the accordion, in order of severity of events that have occurred on them. You can view up to six interfaces at a time. To view interfaces that are not listed, click the **Select Interfaces** drop-down list and select the interface. However, you must clear a previously selected interface to be able to select another interface.

Click the toggle button next to the name of the interface, to view details on link state performance and issues for that interface on a graph. The graph displays two lines showing the high (in orange) and critical (in red) threshold levels related to events that have occurred on the interface. Another line on the graph displays the data points related to events that have occurred on the interface. The color of the data points depends on the color of the interface that you have selected. The color of the interface is assigned automatically. Click any data point to view more information about the corresponding event in a pop-up. The cause for the alert, if any, is also displayed.

To view link state performance and issues for more than one interface on the graph, click the toggle button next to the name of the interface in the Show Link State list. The graphs for the interfaces are displayed one after the other.

You can view link state performance for a week, a day, 3 hours, 1 hour, 30 minutes, or a custom time period. By default, information for the past 30 minutes is displayed. To change this period, click the **Week**, **Day**, **3 hrs**, **1 hr**, **30 mins**, or **Custom** buttons provided above the graph.

The graph auto-refreshes at an interval depending on the time range for which the information is displayed. See [Table 78 on page 306](#) for more information. However, you can refresh the graph at any point by clicking the **Refresh** icon provided above the graph.

You can also click the pop-out button next to the graph to open the graph in a new tab and view all customizations that you made on the graph in the parent tab, in the new tab.

The most critical alert (issues and anomalies) related to link state performance is displayed just above the graph and next to the quick chart. To view other alerts, click the *link* just below the alert. The quick chart displays the link state performance and issues related to the interface that you selected from the

Select Interfaces drop-down list. However, if alerts related to any interface is generated, the interface with the most severe alert is displayed on the quick chart by default.

Alerts, if any, related to link state performance and issues that have occurred are displayed on the graph. You can also open the graph in a new tab. When you open a graph in a new tab, you can view the following information in the new tab as well:

- Alerts related to link state performance.

Alerts are refreshed across all open tabs simultaneously, when:

- An alert is refreshed in any one of the open tabs.
- The last alert fetched was beyond three minutes.
- List of interfaces that you toggled to view from the Select Interfaces drop-down list.

Table 78: Auto-Refresh Rate

Time Range	Auto-Refreshed
Weekly	Every 16.8 hours
Daily	Every 58 minutes
Every three hours	Every 8 minutes
Hourly	Every 3 minutes
Thirty minutes	Every 2 minutes
Custom	No auto-refresh

You can similarly view the graphs and alerts related to port flapping issues.

Software Data and Test Results

Routing Director verifies whether the software (OS) installed on the device is genuine or not. It collects and displays the data listed in the [Table 79 on page 307](#) in the Software accordion of the *Device-Name* page.

You can view the overall reliability of the OS that is installed on the device at the top-right corner of the accordion:

- **Healthy:** The OS installed on the device is genuine, there are no SIRT advisories for the OS and the OS has not reached its end-of-life (EoL).
- **Being Monitored:** The OS is being monitored for issues or vulnerabilities.
- **Action Needed:** The OS has one or more issues, and you must resolve the issues.
- **Urgent Action Needed:** The OS has one or more issues, and you must resolve the issues immediately.



NOTE: For third-party devices, information related to the reliability of the OS is available on the Overview Tab.

Table 79: Fields on the Software Accordion

Field	Description
Vendor	Vendor of the device.
Model	Model of the device; for example, ACX7100-48L. Click the <i>Device-Model</i> link to view the name and IP address of other devices of the same model present in your network, on the Devices with Model <i>Device-Model</i> page.
Licenses	Click the <i>License</i> link to view the list of all active and expected licenses installed on the device. You can view details such as issue date, validity, and expiry date of the licenses installed on the device.
Software	Version of OS installed on the device. Click the <i>Version-Number</i> link to view the device name and IP address of other devices in your network that have the same OS version, on the Devices With Software <i>Version-Number</i> .
End of life - Software	Displays the EoL date for the OS installed on the device.
SIRT Advisories	View the number of Security Incident Response Team (SIRT) advisories that apply to the device and the software that is installed on the device. Click the <i>SIRT advisories count</i> link to view all the SIRT advisories applicable for the device and the OS installed on the device on the Vulnerabilities page (Trust > Vulnerabilities).

Table 79: Fields on the Software Accordion (*Continued*)

Field	Description
Proactive Bug Notifications	
Relevant Events	Lists the latest two events or anomalies in the software installed on the device in order of severity. Hover over View Details to view details of that alert. Click the View Relevant Events link to view the list of all the anomalies present on the device during the past seven days on the Events for <i>Device-Name</i> page.
View Device Documentation	Click to view the documentation for the OS installed on the device.
Upgrade Software	Click to upgrade or downgrade the software that is installed on the device. When you click Upgrade Software, the Upgrade Devices page appears where you can choose the version of the software that you want to upgrade to or downgrade to on the device. To choose the upgrade image, select the device and click the Edit (pencil) icon and select the software version from the Upgrade Image field. Click the check mark to confirm the selection and click OK to start the upgrade or downgrade process.

RELATED DOCUMENTATION

[Upload a Software Image | 366](#)

[Trust and Compliance Overview | 603](#)

Configuration Data and Test Results

Routing Director checks the compliance of the active configuration on the device with the Center for Internet Security (CIS) benchmarks and displays the compliance score. [Table 80 on page 309](#) lists the results of the configuration tests executed on a device and displayed on the Configuration accordion of the *Device-Name* page.

You can view the level of compliance of the configuration committed on the device at the top-right corner of the accordion:

- **Healthy:** The configuration committed is compliant with the CIS benchmarks.
- **Being Monitored:** The configuration committed is being monitored for non-compliance with the CIS benchmarks.
- **Action Needed:** The configuration committed is not as per the CIS benchmarks and hence not compliant. You must intervene to ensure compliance.
- **Urgent Action Needed:** The configuration committed is deviating considerably from CIS benchmarks. You must perform corrective action immediately to ensure compliance.



NOTE: Compliance checks are not supported on third-party devices. As a result, the configuration test results for third-party devices are not available on the Overview Tab.

Table 80: Fields in the Configuration Accordion

Field	Description
Active Version	Active version of the configuration committed on the device. Click the View active config link to view the active configuration on the device.
Other Versions	Backup versions of the configuration. Click a backup configuration link to view the configuration on the Config <i>Device-Name-Date@Time</i> page where, <i>Date</i> and <i>Time</i> are the date and time when the configuration on the device was backed up. Click the Compare link to compare one version of the backup configuration with any other version of the backup configuration on the Config diff - <i>Device-Name</i> page.
Compliance	Displays the most recent compliance score recorded for the active configuration committed on the device. You get a warning if the compliance score of the active configuration is below a specified threshold. Click the Score link to view the compliance scan results and details about the rules that did not meet the criteria defined in the CIS benchmarks document on the Rule Results page (Trust > Compliance).

Table 80: Fields in the Configuration Accordion (*Continued*)

Field	Description
Relevant Events	Lists the latest two alerts and alarms related to the failed compliance checks and failure in committing configuration. Hover over View Details to view details of that alert. Click the View Relevant Events link to view the list of all the events or anomalies present, on the device for the past seven days, on the Events for <i>Device-Name</i> page.
Backup	Click the Backup button to take a backup of the active configuration. A confirmation message appears. Click OK to backup the active configuration.

RELATED DOCUMENTATION

[Trust Score Overview | 621](#)

[Configuration Templates Overview | 375](#)

Routing and MPLS Data and Test Results

SUMMARY

This topic provides information about the results of the tests that Routing Director executes to determine that the states of all BGP, OSPF, IS-IS, RSVP, LSP, and LDP neighbors are healthy.

IN THIS SECTION

- [Overview | 311](#)
- [BGP Routing Details for *Device-Name* Page | 315](#)
- [IS-IS Routing Details for *Device-Name* Page | 319](#)
- [OSPF Routing Details for *Device-Name* Page | 322](#)
- [RSVP Routing Details for *Device-Name* Page | 325](#)
- [LDP Routing Details for *Device-Name* Page | 328](#)

- [LSP Routing Details for *Device-Name*](#)
Page | 330
- [RIB Details | 333](#)
- [FIB Details | 334](#)
- [Traffic Loss | 336](#)
- [Predictor Events | 338](#)

Overview

The Routing and MPLS accordion displays the routing data and results of the tests that Routing Director executes to:

- Determine that the peering states of all BGP, OSPF, IS-IS, RSVP, LSP, and LDP neighbors are healthy and without extensive flaps.



NOTE: A healthy peering state is marked **Up/Active**. An unhealthy peering state is marked **Down**.

- Validate that the expected numbers of entries are available in the routing and forwarding tables.
- View alerts related to traffic loss (**Traffic Loss**) and exceptions on a device that may have caused traffic loss (**Predictor Events**).

To access the Routing and MPLS accordion, navigate to the **Observability > Health > Troubleshoot Devices** page. Click a device name to access the *Device-Name* page. Click the **Routing and MPLS (accordion)** in the Overview tab.

The top-right corner of the accordion displays state of the overall health of the routing components. The various states are:

- **Healthy**—The BGP, OSPF, IS-IS, RSVP, LSP, and LDP neighbors are healthy. The entries in the routing and forwarding tables are accurate.
- **Being Monitored**—The health of the routing neighbors and entries in the routing and forwarding tables is being monitored.
- **Action Needed**—The routing neighbors and entries in the routing and forwarding tables have issues that you must address.
- **Urgent Action Needed**—The routing neighbors and entries in the routing and forwarding tables have issues that you must address immediately.

Table 81 on page 312 lists the results of the routing tests.



NOTE: Routing tests are not supported on third-party devices. As a result, the routing data and test results for third-party devices are not available on the Overview Tab.

Table 81: Results of Routing Tests

Field	Description
BGP	Displays the total number of Border Gateway Protocol (BGP) neighbors available, and the number of unhealthy BGP neighbors. Unhealthy BGP neighbors, if any, are also displayed. Click the link next to BGP to view information about BGP peering status, BGP flaps, routes that are advertised to a neighbor, routes received from a neighbor, and alerts. NOTE: A healthy peering state is marked Up/Active. An unhealthy peering state is marked Down. See "BGP Routing Details for Device-Name Page" on page 315 for more information.
ISIS	Displays the total number of Intermediate System-to-Intermediate System (IS-IS) neighbors available, and the number of unhealthy IS-IS neighbors. Unhealthy IS-IS neighbors, if any, are also displayed. Click the link next to ISIS to view information about IS-IS adjacency state, IS-IS flaps, IS-IS drops, and alerts. See "IS-IS Routing Details for Device-Name Page" on page 319 for more information.
OSPF	Displays the total number of Open Shortest Path First (OSPF) neighbors available, and the number of unhealthy OSPF neighbors. Unhealthy OSPF neighbors, if any, are also displayed. Click the link next to OSPF to view information about OSPF interface state, OSPF hello protocol, OSPF input/output (I/O) errors, and alerts. See "OSPF Routing Details for Device-Name Page" on page 322 for more information.

Table 81: Results of Routing Tests *(Continued)*

Field	Description
RSVP	Displays the total number of Resource Reservation Protocol (RSVP) neighbors available, and the number of unhealthy RSVP neighbors. Unhealthy RSVP neighbors, if any, are also displayed. Click the link next to RSVP to view information about the state of RSVP neighbors, RSVP global traffic engineering (TE) errors, TE interface errors, and alerts. See "RSVP Routing Details for Device-Name Page" on page 325 for more information.
LDP	Displays the total number of Label Distribution Protocol (LDP) neighbors available, and the number of unhealthy LDP neighbors. Unhealthy LDP neighbors, if any, are also displayed. Click the link next to LDP to view information about the state of LDP peers, and alerts. See "LDP Routing Details for Device-Name Page" on page 328 for more information.
LSP	Displays the total number of times an Label-Switched Path (LSP) flaps, and the LSP state (Up or Down). Unhealthy LSP peers, if any, are also displayed. Click the link next to LSP to view information about the LSP peer state, LSP flaps, and alerts. See "LSP Routing Details for Device-Name Page" on page 330 for more information.
RIB	Displays the total number of routes in the routing information base (RIB), also known as routing table. Click the icon next to RIB to view a pop-up of a list of routing tables in the RIB, and the total number of routes in each table. Click the link next to RIB to view detailed information about RIB, learned routes, total number of active routes, and alerts. See "RIB Details" on page 333 for more information.

Table 81: Results of Routing Tests *(Continued)*

Field	Description
FIB	Displays the total number of routes in the forwarding information base (FIB), also known as forwarding table. Click the icon next to FIB to view a pop-up of a list of forwarding tables in the FIB, and the total number of routes in each table. Click the link next to FIB to view detailed information about FIB, learned routes, total number of active routes, and alerts. See "FIB Details" on page 334 for more information.
Predictor Events	Displays the number of alerts caused due to routing, forwarding, and OS exceptions on the device. Click the Alerts link to view the detailed information about predictor events. You can also view other events (correlated events) that may have caused the predictor events. See "Predictor Events" on page 338 for details.
Traffic Loss	Displays the number of events indicating packet drops on the device. Click the Alerts link to view the details of the packet drop alerts on the Traffic Loss page. The traffic loss page displays two graphs: • A line graph indicating packet drops at different times • A line graph indicating the packet input and output rates and the drop rate, all on the same graph. Clicking an alert icon, filters and displays the alert on the Alerts table. See "Traffic Loss" on page 336 for details.

Table 81: Results of Routing Tests (*Continued*)

Field	Description
Relevant Events	Displays two issues or anomalies with respect to the routing events in order of severity. Hover over View Details to view more information about an issue in a pop-up. Click View All Relevant Events to view all routing events or anomalies, present on the device, on the Events for <i>Device-Name</i> page. You can view relevant events for the past seven days. NOTE: The Relevant Events section displays alerts related to bgp-neighbor-flap, bgp-neighbor-status, advertised-routes, received-routes, isis-adjacency-flap, isis-adjacency-state, isis-csnp-drops, isis-esh-drops, isis-iih-drops, isis-ish-drops, isis-lsp-drops, isis-psnp-drops, isis-unknown-drops, ospf-hello-received, ospf-hello-sent, ospf-neighbor-state, ospf-io-errors, lsp-flaps, lsp-state, ldp-session-state, rsvp-neighbor-state, rsvp-te-global-errors, rsvp-te-interface-errors, route-table-total-routes-count, protocol-routes-count, fib-route-count, lldp-session-state, lfm-interface-discovery-state, evpn-mac-count, evpn-instance-state, evpn-neighbor-state, evpn-pe-interface-state, evpn-state, l3vpn-pe-interface-state, l3vpn-state, l3vpn-pe-interface-state, l3vpn-state, and l3vpn-pe-interface-state.



NOTE: KPI-related data may not be displayed on the accordions and on the graphs when:

- There are device errors
- The data displayed is derived from or is dependent on another KPI that is not being streamed

BGP Routing Details for *Device-Name* Page

To access the BGP routing details on the Routing Director GUI, click **Observability > Health > Troubleshoot Devices > Device-Name > Overview > Routing and MPLS (accordion) > BGP**. Click the *link* next to each BGP component to view its details.

The four accordions on this page provide information about BGP peering status, BGP flaps, routes that are advertised to a neighbor, and routes received from a neighbor. [Table 82 on page 316](#) describes the accordions.

Table 82: Accordions on the BGP Routing Details for *Device-Name* Page

Accordion	Description
Up/Down	View the BGP peering status and alerts for up to six peers on a graph at a time. These peers are listed in order of the severity of the alerts that have occurred on them. The BGP peer with the most critical alerts appears at the top of the list. NOTE: A healthy peering state is marked Up/Active. An unhealthy peering state is marked Down. Click the toggle button next to a BGP peer in the Show Peers list to view the performance and alerts of that peer in a graph. See "Performance Graphs" on page 317 for more information.
Advertised Routes	View information and alerts on BGP advertised routes for up to six routes on a graph at a time. The routes are listed in the order of severity of the alerts that have occurred on them. The BGP route with the most critical alerts appears at the top of the list. Click the toggle button next to a BGP route in the Show Routes list to view the performance and alerts of that route in a graph. See "Performance Graphs" on page 317 for more information.
Received Routes	View information and alerts on BGP received routes for up to six routes on a graph at a time. The routes are listed in the order of severity of the alerts that have occurred on them. The BGP route with the most critical alerts appears at the top of the list. Click the toggle button next to a BGP route in the Show Routes list to view the performance and alerts of that route in a graph. See "Performance Graphs" on page 317 for more information.
BGP Flaps	View BGP flapping information and alerts for up to six peers on a graph at a time. These peers are listed in order of the severity of the alerts that have occurred on them. The BGP peers with the most critical alerts appear at the top of the list. Click the toggle button next to a BGP peer in the Show Peers list to view the performance and alerts of that peer in a graph. See "Performance Graphs" on page 317 for more information.

Performance Graphs

The graphs on the BGP Details for *Device-Name* page display information about BGP peering status, BGP flaps, routes that are advertised to a neighbor, and routes received from a neighbor. You can also view information on alerts and breaches, if any, on these graphs. [Figure 25 on page 317](#) shows details about the BGP advertised routes on the graph.

Figure 25: Advertised Routes Accordion



The routes that are advertised are listed on the left of the Advertised Routes accordion, in order of the severity of alerts that have occurred on them. You can view up to six routes at a time. To view routes that are not listed, click the **Show Routes** drop-down list and select the route. However, you must clear a previously selected route to be able to select another route.

Click the toggle button next to the name of the BGP route to view the performance of that route in a graph. The graph displays two lines showing the high (in orange) and critical (in red) threshold levels related to the alerts that have occurred. Another line on the graph displays the data points that are related to alerts that have occurred on the route. The color of the data points depends on the color of the route that you have selected. The color of the route is assigned automatically. Click any data point to view more information about the corresponding alert in a pop-up. The cause for the alert, if any, is also displayed.

To view the performance of more than one route on the graph, click the toggle button next to the name of the route in the Show Routes list. Details of the routes displayed on the graph are also listed in a table below the graph. You can also click the option buttons on the left of a route name in the table to highlight the graph for that route.

You can view the advertised route information for a week, a day, 3 hours, 1 hour, 30 minutes, or a custom time period. By default, information for the past 30 minutes is displayed. To change this period, click the **Week**, **Day**, **3 hrs**, **1 hr**, **30 mins**, or **Custom** buttons provided above the graph.

The graph auto-refreshes at an interval depending on the time range for which the information is displayed. See [Table 83 on page 318](#) for more information. However, you can refresh the graph at any point by clicking the refresh icon provided above the graph.

Table 83: Auto-Refresh Rate

Time Range	Auto-Refreshed
Weekly	Every 16.8 hours
Daily	Every 58 minutes
Every three hours	Every 8 minutes
Hourly	Every 3 minutes
Thirty minutes	Every 2 minutes
Custom	No auto-refresh

You can also click the pop-out button next to the graph to open the graph in a new tab and view all customizations that you made on the graph in the parent tab, in the new tab.

The most critical alert (issues and anomalies) is displayed just above the graph and next to the quick chart. To view other alerts, click the [link](#) just below the alert. The quick chart displays the performance of the route that you selected from the Select Routes drop-down list. However, if alerts related to the performance of any route is generated, the route with the most severe alert is displayed on the quick chart by default.

Alerts, if any, related to advertised routes are displayed on the graph, and in the table below the graph. You can also open the graph in a new tab by clicking the pop-out button. When you open a graph in a new tab, you can view the following information in the new tab as well:

- Alerts that are present in the home tab are also present in the newly opened tab as well.
- Alerts are refreshed across all open tabs simultaneously.
- List of routes that you toggled to view from the Select Routes drop-down list is available across all tabs.
- Route that you selected from the table below the graph are available across all tabs.

You can similarly view the graphs and alerts related to BGP peering status, routes received from a neighbor, and BGP flapping information.

IS-IS Routing Details for *Device-Name* Page

To access the IS-IS routing details on the Routing Director GUI, click **Observability > Health > Troubleshoot Devices > *Device-Name* > Overview > Routing and MPLS (accordion) > ISIS**. Click the *link* next to each hardware component to view its details.

The three accordions on this page provide information about IS-IS adjacency state, IS-IS flaps, and IS-IS drops. [Table 84 on page 319](#) describes the accordions.

Table 84: Accordion on the IS-IS Routing Details for *Device-Name* Page

Accordion	Description
ISIS Adjacency State	View information, alerts, and state of IS-IS adjacencies for up to six interfaces on a graph at a time. These interfaces are listed in the order of severity of the alerts that have occurred on them. The interfaces with the most critical alerts appear at the top of the list. Click the toggle button next to an interface in the Show Interfaces list to view the performance and alerts of that interface in a graph. See "Performance Graphs" on page 320 for more information.
Adjacency Flap	View IS-IS flapping information and alerts for up to six interfaces on a graph at a time. These interfaces are listed in the order of severity of the alerts that have occurred on them. The interfaces with the most critical alerts appear at the top of the list. Click the toggle button next to an interface in the Show Interfaces list to view the performance and alerts of that interface in a graph. See "Performance Graphs" on page 320 for more information.
ISIS Drops	View information and alerts related to IS-IS drops for up to six interfaces on a graph at a time. These interfaces are listed in the order of severity of the alerts that have occurred on them. The interfaces with the most critical alerts appear at the top of the list. Click the toggle button next to an interface in the Show Interfaces list to view the performance and alerts of that interface in a graph. See "Performance Graphs" on page 320 for more information.

Performance Graphs

The graphs on the ISIS Details for *Device-Name* page display information about IS-IS adjacency state, IS-IS flaps, and IS-IS drops. You can also view information on alerts and breaches, if any, on these graphs.

[Figure 26 on page 320](#) shows details about the state of IS-IS adjacencies on the graph.

Figure 26: ISIS Adjacency State Accordion



The interfaces are listed on the left of the ISIS Adjacency State accordion, in order of the severity of alerts that have occurred on them. You can view up to six interfaces at a time. To view interfaces that are not listed, click the **Select Interfaces** drop-down list and select the interface. However, you must clear a previously selected interface to be able to select another interface.

Click the toggle button next to the name of the interface to view the performance of that interface in a graph. The graph displays two lines showing the high (in orange) and critical (in red) threshold levels related to the alerts that have occurred. Another line on the graph displays the data points that are related to alerts that have occurred on the interface. The color of the data points depends on the color of the interface that you have selected. The color of the interface is assigned automatically. Click any data point to view more information about the corresponding event in a pop-up. The cause for the alert, if any, is also displayed.

To view the performance of more than one interface on the graph, click the toggle button next to the name of the interface in the Show Interfaces list. Details of the interfaces displayed on the graph are also listed in a table below the graph. You can also click the option buttons on the left of an interface name in the table to highlight the graph for that interface.

You can view the adjacency performance for a week, a day, 3 hours, 1 hour, 30 minutes, or a custom time period. By default, information for the past 30 minutes is displayed. To change this period, click the **Week**, **Day**, **3 hrs**, **1 hr**, **30 mins**, or **Custom** buttons provided above the graph.

The graph auto-refreshes at an interval depending on the time range for which the information is displayed. See [Table 85 on page 321](#) for more information. However, you can refresh the graph at any point by clicking the refresh icon provided above the graph.

Table 85: Auto-Refresh Rate

Time Range	Auto-Refreshed
Weekly	Every 16.8 hours
Daily	Every 58 minutes
Every three hours	Every 8 minutes
Hourly	Every 3 minutes
Thirty minutes	Every 2 minutes
Custom	No auto-refresh

You can also click the pop-out button next to the graph to open the graph in a new tab and view all customizations that you made on the graph in the parent tab, in the new tab.

The most critical alert (issues and anomalies) is displayed just above the graph and next to the quick chart. To view other alerts, click the *link* just below the alert. The quick chart displays the performance of the interface that you selected from the Select Interfaces drop-down list. However, if alerts related to the performance of any interface is generated, the interface with the most severe alert is displayed on the quick chart by default.

Alerts, if any, related to interfaces are displayed on the graph, and in the table below the graph. You can also open the graph in a new tab by clicking the pop-out button. When you open a graph in a new tab, you can view the following information in the new tab as well:

- Alerts that are present in the home tab are also present in the newly opened tab as well.
- Alerts are refreshed across all open tabs simultaneously.
- List of interfaces that you toggled to view from the Select Interfaces drop-down list is available across all tabs.
- Interface that you selected from the table below the graph are available across all tabs.

You can similarly view the graphs and alerts related to IS-IS flaps and IS-IS drops.

OSPF Routing Details for *Device-Name* Page

To access the OSPF routing details on the Routing Director GUI, click **Observability > Health > Troubleshoot Devices > *Device-Name* > Overview > Routing and MPLS (accordion) > OSPF**. Click the *link* next to each OSPF component to view its details.

The four accordions on this page provide information on the OSPF interface state, OSPF hello protocol, and OSPF input/output (I/O) errors. [Table 86 on page 322](#) describes the accordions.

Table 86: Accordions on the OSPF Routing Details for *Device-Name* Page

Accordion	Description
OSPF IPv6 State	View information, alerts, and OSPFv3 interface state for IPv6 on up to six interfaces on a graph at a time. These interfaces are listed in the order of severity of the alerts that have occurred on them. The interfaces with the most critical alerts appear at the top of the list. Click the toggle button next to an interface in the Show Interfaces list to view the performance and alerts of that interface in a graph. See "Performance Graphs" on page 323 for more information.
OSPF State	View information, alerts, and OSPF interface state for up to six interfaces on a graph at a time. These interfaces are listed in the order of severity of the alerts that have occurred on them. The interfaces with the most critical alerts appear at the top of the list. Click the toggle button next to an interface in the Show Interfaces list to view the performance and alerts of that interface in a graph. See "Performance Graphs" on page 323 for more information.
OSPF Hello Received	View information and alerts related to OSPF Hello protocol (packets received) on a graph. See "Performance Graphs" on page 323 for more information.
OSPF Hello Sent	View information and alerts related to OSPF Hello protocol (packets sent) on a graph. See "Performance Graphs" on page 323 for more information.
OSPF I/O Errors	View information and alerts related to OSPF input and output errors on a graph. See "Performance Graphs" on page 323 for more information.

Performance Graphs

The graphs on the OSPF Details for *Device-Name* page display information about OSPF interface state, OSPF hello protocol, and OSPF input/output (I/O) errors. You can also view information on alerts and breaches, if any, on these graphs. [Figure 27 on page 323](#) shows details about the OSPF interface state on the graph.

Figure 27: OSPF State Accordion



The interfaces are listed on the left of the OSPF State accordion, in order of the severity of alerts that have occurred on them. You can view up to six interfaces at a time. To view interfaces that are not listed, click the **Select Interfaces** drop-down list and select the interface. However, you must clear a previously selected interface to be able to select another interface.



NOTE: The OSPF Hello Received, OSPF Hello Sent, and OSPF I/O Errors accordions in the OSPF Routing Details for *Device-Name* page do not have interfaces listed next to the graph.

Click the toggle button next to the name of the interface to view the performance of that interface in a graph. The graph displays two lines showing the high (in orange) and critical (in red) threshold levels related to the alerts that have occurred. Another line on the graph displays the data points related to alerts that have occurred on the interface. The color of the data points depends on the color of the interface that you have selected. The color of the interface is assigned automatically. Click any data point to view more information about the corresponding event in a pop-up. The cause for the alert, if any, is also displayed.

To view the performance of more than one interface on the graph, click the toggle button next to the name of the interface in the Show Interfaces list.

You can view the performance for a week, a day, 3 hours, 1 hour, 30 minutes, or a custom time period. By default, information for the past 30 minutes is displayed. To change this period, click the **Week**, **Day**, **3 hrs**, **1 hr**, **30 mins**, or **Custom** buttons provided above the graph.

The graph auto-refreshes at an interval depending on the time range for which the information is displayed. See [Table 87 on page 324](#) for more information. However, you can refresh the graph at any point by clicking the refresh icon provided above the graph.

Table 87: Auto-Refresh State

Time Range	Auto-Refreshed
Weekly	Every 16.8 hours
Daily	Every 58 minutes
Every three hours	Every 8 minutes
Hourly	Every 3 minutes
Thirty minutes	Every 2 minutes
Custom	No auto-refresh

You can also click the pop-out button next to the graph to open the graph in a new tab and view all customizations that you made on the graph in the parent tab, in the new tab. When you open a graph in a new tab, you can view the following information in the new tab as well:

The most critical alert (issues and anomalies) is displayed just above the graph and next to the quick chart. To view other alerts, click the *link* just below the alert. The quick chart displays the performance of the interface that you selected from the Select Interfaces drop-down list. However, if alerts related to the performance of any interface is generated, the interface with the most severe alert is displayed on the quick chart by default.

Alerts, if any, related to interfaces are displayed on the graph, and in the table below the graph. You can also open the graph in a new tab by clicking the pop-out button. When you open a graph in a new tab, you can view the following information in the new tab as well:

- Alerts that are present in the home tab are also present in the newly opened tab as well.
- Alerts are refreshed across all open tabs simultaneously.
- List of interfaces that you toggled to view from the Select Interfaces drop-down list is available across all tabs.
- Interface that you selected from the table below the graph are available across all tabs.

You can similarly view the graphs and alerts related to OSPF hello protocol, and OSPF input/output (I/O) errors.

RSVP Routing Details for *Device-Name* Page

To access the RSVP routing details on the Routing Director GUI, click **Observability > Health > Troubleshoot Devices > *Device-Name* > Overview > Routing and MPLS (accordion) > RSVP**. Click the *link* next to each RSVP component to view its details.

The four accordions on this page provide information on the RSVP neighbors, RSVP global traffic engineering (TE) errors, and TE interface errors. [Table 88 on page 325](#) describes the accordions.

Table 88: Accordions on the RSVP Routing Details for *Device-Name* Page

Accordion	Description
RSVP Neighbor State	View information and alerts related to the state of RSVP peers for up to six peers on a graph at a time. These peers are listed in order of the severity of the alerts that have occurred on them. The peers with the most critical alerts appear at the top of the list. Click the toggle button next to a peer in the Show Peers list to view the performance and alerts of that peer in a graph. See "Performance Graphs" on page 326 for more information.
RSVP TE Global Errors	View information and alerts related to TE errors for up to six routing instances on a graph at a time. These instances are listed in the order of severity of the alerts that have occurred on them. The instances with the most critical alerts appear at the top of the list. Click the toggle button next to an instance in the Show Instances list to view the performance and alerts of that instance in a graph. See "Performance Graphs" on page 326 for more information.
RSVP TE Interface Errors	View information and alerts related to TE errors for up to six interfaces on a graph at a time. These interfaces are listed in the order of severity of the alerts that have occurred on them. The interfaces with the most critical alerts appear at the top of the list. Click the toggle button next to an interface in the Show Interfaces list to view the performance and alerts of that interface in a graph. See "Performance Graphs" on page 326 for more information.

Performance Graphs

The graphs on the RSVP Details for *Device-Name* page display information about RSVP neighbors, RSVP global TE errors, and TE interface errors. You can also view information on alerts and breaches, if any, on these graphs. [Figure 28 on page 326](#) shows details about the RSVP TE interface errors on the graph.

Figure 28: RSVP TE Interface Accordion



The interfaces are listed on the left of the RSVP TE Interface accordion, in the order of severity of alerts that have occurred on them. The interfaces that are listed are based on the errors that you select from the drop-down list provided below Show Interfaces. Click the drop-down, and select an error from the list to view interfaces with that error.

You can view up to six interfaces at a time. To view interfaces that are not listed, click the **Select Interfaces** drop-down list and select the interface. However, you must clear a previously selected interface to be able to select another interface.

Click the toggle button next to the name of the interface to view the performance of that interface in a graph. The graph displays two lines showing the high (in orange) and critical (in red) threshold levels related to the alerts that have occurred. Another line on the graph displays the data points related to alerts that have occurred on the interface. The color of the data points depends on the color of the interface that you have selected. The color of the interface is assigned automatically. Click any data point to view more information about the corresponding event in a pop-up. The cause for the alert, if any, is also displayed.

To view the performance of more than one interface on the graph, click the toggle button next to the name of the interface in the Show Interfaces list. Details of the interfaces displayed on the graph are also listed in a table below the graph. You can also click the option buttons on the left of an interface name in the table to highlight the graph for that interface.

You can view the performance for a week, a day, 3 hours, 1 hour, 30 minutes, or a custom time period. By default, information for the past 30 minutes is displayed. To change this period, click the **Week**, **Day**, **3 hrs**, **1 hr**, **30 mins**, or **Custom** buttons provided above the graph.

The graph auto-refreshes at an interval depending on the time range for which the information is displayed. See [Table 89 on page 327](#) for more information. However, you can refresh the graph at any point by clicking the refresh icon provided above the graph.

Table 89: Auto-Refresh State

Time Range	Auto-Refreshed
Weekly	Every 16.8 hours
Daily	Every 58 minutes
Every three hours	Every 8 minutes
Hourly	Every 3 minutes
Thirty minutes	Every 2 minutes
Custom	No auto-refresh

You can also click the pop-out button next to the graph to open the graph in a new tab and view all customizations that you made on the graph in the parent tab, in the new tab.

The most critical alert (issues and anomalies) is displayed just above the graph and next to the quick chart. To view other alerts, click the [link](#) just below the alert. The quick chart displays the performance of the interface that you selected from the Select Interfaces drop-down list. However, if alerts related to the performance of any interface is generated, the interface with the most severe alert is displayed on the quick chart by default.

Alerts, if any, related to interfaces are displayed on the graph, and in the table below the graph. You can also open the graph in a new tab by clicking the pop-out button. When you open a graph in a new tab, you can view the following information in the new tab as well:

- Alerts that are present in the home tab are also present in the newly opened tab as well.
- Alerts are refreshed across all open tabs simultaneously.
- List of interfaces that you toggled to view from the Select Interfaces drop-down list is available across all tabs.
- Interface that you selected from the table below the graph are available across all tabs.

You can similarly view the graphs and alerts related to RSVP neighbors and RSVP global traffic engineering TE errors.

LDP Routing Details for *Device-Name* Page

To access the LDP routing details on the Routing Director GUI, click **Observability > Health > Troubleshoot Devices > *Device-Name* > Overview > Routing and MPLS (accordion) > LDP**. Click the *link* next to each LDP component to view its details.

The accordion on this page provides information on the state of LDP peers. [Table 90 on page 328](#) describes the accordion.

Table 90: Accordions on the LDP Routing Details for *Device-Name* Page

Accordion	Description
LDP Neighbor State	View information on the state of LDP peers. These peers are listed in order of the severity of the alerts that have occurred on them. The peers with the most critical alerts appear at the top of the list. Click the toggle button next to a peer in the Show Peers list to view the performance and alerts of that peer in a graph. See "Performance Graphs" on page 328 for more information.

Performance Graphs

The graphs on the LDP Details for *Device-Name* page display information about the state of LDP peers. You can also view information on alerts and breaches, if any, on these graphs. For more information, see [Figure 29 on page 329](#).

Figure 29: LDP Neighbor State Accordion



The peers are listed on the left of the LDP Neighbor State accordion, in order of the severity of alerts that have occurred on them. You can view up to six peers at a time. To view peers that are not listed, click the **Select Peers** drop-down list and select the peer. However, you must clear a previously selected peer to be able to select another peer.

Click the toggle button next to the name of the peer to view the performance of that peer in a graph. The graph displays two lines showing the high (in orange) and critical (in red) threshold levels related to the alerts that have occurred. Another line on the graph displays the data points related to alerts that have occurred on the peer. The color of the data points depends on the color of the peer that you have selected. The color of the peer is assigned automatically. Click any data point to view more information about the corresponding event in a pop-up. The cause for the alert, if any, is also displayed.

To view the performances of more than one peer on the graph, click the toggle button next to the name of the peer in the Show Peers list. Details of the interfaces displayed on the graph are also listed in a table below the graph. You can also click the option buttons on the left of a peer name in the table to highlight the graph for that peer.

You can view performance for a week, a day, 3 hours, 1 hour, 30 minutes, or a custom time period. By default, information for the past 30 minutes is displayed. To change this period, click the **Week**, **Day**, **3 hrs**, **1 hr**, **30 mins**, or **Custom** buttons provided above the graph.

The graph auto-refreshes at an interval depending on the time range for which the information is displayed. See [Table 91 on page 330](#) for more information. However, you can refresh the graph at any point by clicking the refresh icon provided above the graph.

Table 91: Auto-Refresh State

Time Range	Auto-Refreshed
Weekly	Every 16.8 hours
Daily	Every 58 minutes
Every three hours	Every 8 minutes
Hourly	Every 3 minutes
Thirty minutes	Every 2 minutes
Custom	No auto-refresh

You can also click the pop-out button next to the graph to open the graph in a new tab and view all customizations that you made on the graph in the parent tab, in the new tab.

The most critical alert (issues and anomalies) is displayed just above the graph and next to the quick chart. To view other alerts, click the *link* just below the alert. The quick chart displays the performance of the peer that you selected from the Select Peers drop-down list. However, if alerts related to the performance of any peer is generated, the peer with the most severe alert is displayed on the quick chart by default.

Alerts, if any, related to peers are displayed on the graph, and in the table below the graph. You can also open the graph in a new tab by clicking the pop-out button. When you open a graph in a new tab, you can view the following information in the new tab as well:

- Alerts that are present in the home tab are also present in the newly opened tab as well.
- Alerts are refreshed across all open tabs simultaneously.
- List of peers that you toggled to view from the Select Interfaces drop-down list is available across all tabs.
- Peer that you selected from the table below the graph are available across all tabs.

LSP Routing Details for *Device-Name* Page

To access the LSP routing details on the Routing Director GUI, click **Observability > Health > Troubleshoot Devices > *Device-Name* > Overview > Routing and MPLS (accordion) > LSP**. Click the *link* next to each LSP component to view its details.

The two accordions on this page provide information on the state of LSP peers and LSP flaps. [Table 92 on page 331](#) describes the accordion.

Table 92: Accordions on the LSP Routing Details for *Device-Name* Page

Accordion	Description
LSP Flap	View LSP flapping issues for various LSP peers. These peers are listed in order of the severity of the alerts that have occurred on them. The peers with the most critical alerts appear at the top of the list. Click the toggle button next to a peer in the Show Peers list to view the performance and alerts of that peer in a graph. See "Performance Graphs" on page 331 for more information.
LSP State	View the state of LSP peers. These peers are listed in order of the severity of the alerts that have occurred on them. The peers with the most critical alerts appear at the top of the list. Click the toggle button next to a peer in the Show Peers list to view the performance and alerts of that peer in a graph. See "Performance Graphs" on page 331 for more information.

Performance Graphs

The graphs on the LSP Details for *Device-Name* page display information about the state of LSP peers and LSP flaps. You can also view information on alerts and breaches, if any, on these graphs.

As an example, consider the LSP State accordion on the LSP Details for *Device-Name* page. The peers are listed on the left of the accordion, in the order of severity of alerts that have occurred on them. You can view up to six peers at a time. To view peers that are not listed, click the **Select Peers** drop-down list and select the peer. However, you must clear a previously selected peer to be able to select another peer.

Click the toggle button next to the name of the peer to view the performance of that peer in a graph. The graph displays two lines showing the high (in orange) and critical (in red) threshold levels related to the alerts that have occurred. Another line on the graph displays the data points related to alerts that have occurred on the peer. The color of the data points depends on the color of the peer that you have selected. The color of the peer is assigned automatically. Click any data point to view more information about the corresponding event in a pop-up. The cause for the alert, if any, is also displayed.

To view the performance of more than one peer on the graph, click the toggle button next to the name of the peer in the Show Peers list. Details of the interfaces displayed on the graph are also listed in a

table below the graph. You can also click the option buttons on the left of a peer name in the table to highlight the graph for that peer.

You can view performances for a week, a day, 3 hours, 1 hour, 30 minutes, or a custom time period. By default, information for the past 30 minutes is displayed. To change this period, click the **Week**, **Day**, **3 hrs**, **1 hr**, **30 mins**, or **Custom** buttons provided above the graph.

The graph auto-refreshes at an interval depending on the time range for which the information is displayed. See [Table 93 on page 332](#) for more information. However, you can refresh the graph at any point by clicking the refresh icon provided above the graph.

Table 93: Auto-Refresh State

Time Range	Auto-Refreshed
Weekly	Every 16.8 hours
Daily	Every 58 minutes
Every three hours	Every 8 minutes
Hourly	Every 3 minutes
Thirty minutes	Every 2 minutes
Custom	No auto-refresh

You can also click the pop-out button next to the graph to open the graph in a new tab and view all customizations that you made on the graph in the parent tab, in the new tab.

The most critical alert (issues and anomalies) is displayed just above the graph and next to the quick chart. To view other alerts, click the *link* just below the alert. The quick chart displays the performance of the peer that you selected from the Select Peers drop-down list. However, if alerts related to the performance of any peer is generated, the peer with the most severe alert is displayed on the quick chart by default.

Alerts, if any, related to peers are displayed on the graph, and in the table below the graph. You can also open the graph in a new tab by clicking the pop-out button. When you open a graph in a new tab, you can view the following information in the new tab as well:

- Alerts that are present in the home tab are also present in the newly opened tab as well.
- Alerts are refreshed across all open tabs simultaneously.

- List of peers that you toggled to view from the Select Interfaces drop-down list is available across all tabs.
- Peer that you selected from the table below the graph are available across all tabs.

You can similarly view the graphs and alerts related to LSP flaps on this page.

RIB Details

To access the RIB details on the Routing Director GUI, click **Observability > Health > Troubleshoot Devices > *Device-Name* > Overview > Routing and MPLS (accordion) > RIB *data-link***.

The RIB Details page consists of the Route Count and Learned Route sections. See [Figure 30 on page 333](#).

Figure 30: RIB Details Page



The Route Count section displays the total routes available and its performance over a period of time. The routes that are displayed depend on the routing table that you select. In [Figure 30 on page 333](#), the `inet.0` table is selected. To view routes that belong to other routing tables, click the drop-down in the top-right corner of the page and select another table.

The Learned Routes section of the RIB Details page is updated based on the routing table that you selected in the Route Count section. Click the arrow next to `inet.0` to view information related to the routes in the table. The total routes and the total active routes for each route type (direct, local, and static routes) are displayed. Further, click **Add to Chart** to view this route on the graph in the Route Count section. To remove or reset the graph, click **Remove from Chart** or **Reset Charts**, respectively.

You can view performance of a route for a week, a day, 3 hours, 1 hour, 30 minutes, or a custom time period. By default, information for the past 30 minutes is displayed. To change this period, click the **Week**, **Day**, **3 hrs**, **1 hr**, **30 mins**, or **Custom** buttons provided above the graph.

The graph auto-refreshes at an interval depending on the time range for which the information is displayed. See [Table 94 on page 334](#) for more information. However, you can refresh the graph at any point by clicking the refresh icon provided above the graph.

Table 94: Auto-Refresh State

Time Range	Auto-Refreshed
Weekly	Every 16.8 hours
Daily	Every 58 minutes
Every three hours	Every 8 minutes
Hourly	Every 3 minutes
Thirty minutes	Every 2 minutes
Custom	No auto-refresh

You can also click the pop-out button next to the graph to open the graph in a new tab and view all customizations that you made on the graph in the parent tab, in the new tab.

FIB Details

To access the FIB details on the Routing Director GUI, click **Observability > Health > Troubleshoot Devices > *Device-Name* > Overview > Routing and MPLS (accordion) > FIB**. Click the *link* next to each FIB component to view its details.

The FIB Details page consists of the Route Count and Learned Route sections. See [Figure 31 on page 335](#).

Figure 31: FIB Details Page



The Route Count section displays the total routes available and its performance over a period of time. The routes that are displayed depend on the forwarding table that you select. In [Figure 31 on page 335](#), the `default.iso` table is selected. To view routes that belong to other forwarding tables, click the drop-down in the top-right corner of the page and select another table.

The Learned Routes section of the FIB Details page is updated based on the forwarding table that you selected in the Route Count section. Click the arrow next to `default.iso` to view information related to the routes in the table. The total routes and the total active routes for each route type (direct, local, and static routes) are displayed. Further, click **Add to Chart** to view this route on the graph in the Route Count section. To remove or reset the graph, click **Remove from Chart** or **Reset Charts**, respectively.

You can view performance of a route for a week, a day, 3 hours, 1 hour, 30 minutes, or a custom time period. By default, information for the past 30 minutes is displayed. To change this period, click the **Week**, **Day**, **3 hrs**, **1 hr**, **30 mins**, or **Custom** buttons provided above the graph.

The graph auto-refreshes at an interval depending on the time range for which the information is displayed. See [Table 95 on page 335](#) for more information. However, you can refresh the graph at any point by clicking the refresh icon provided above the graph.

Table 95: Auto-Refresh State

Time Range	Auto-Refreshed
Weekly	Every 16.8 hours
Daily	Every 58 minutes

Table 95: Auto-Refresh State *(Continued)*

Time Range	Auto-Refreshed
Every three hours	Every 8 minutes
Hourly	Every 3 minutes
Thirty minutes	Every 2 minutes
Custom	No auto-refresh

You can also click the pop-out button next to the graph to open the graph in a new tab and view all customizations that you made on the graph in the parent tab, in the new tab.

Traffic Loss



NOTE: Traffic Loss is a beta feature in this release.

The Traffic Loss page displays details about the traffic loss that has occurred on the device. To view the Traffic Loss page:

- For a device in service, click **Observability > Health > Troubleshoot Devices > *Device-Name* > Overview > Routing and MPLS accordion** > click Traffic Loss Alert link.
- For a device during onboarding, click **Inventory > Device Onboarding > Onboarding Dashboard > *Device-Name* > Routing and MPLS accordion** > click Traffic Loss Alert link.

Figure on page 337 shows the Traffic Loss page. On the Traffic Loss page, you can see:

- A graph showing packet drop rate.
- A graph showing packet input rate, packet output rate, and packet drop rate.

Click a KPI to turn off the graph for that KPI. For example, clicking **Input Rate** turns off the graph of input packet rate.

- An alerts table listing the alerts raised on the device. You can see the following alerts:
 - A red alert icon (critical) indicates that blackhole is detected
 - An orange alert icon (major) indicates that packet drop anomalies are present but blackhole is not detected.

- A normal small circle icon indicates that blackhole is not detected and packet drop anomalies are not present.

Clicking an alert icon on the graph filters the corresponding alert in the Alerts table.

Figure 32: Traffic Loss page



You can filter data in the Alerts table by the severity, alert name, exception, exception type, and the PFE from which the alert is raised. Select an alert and click **View details** to view details on the Traffic Loss Details page. The Traffic Loss Details page lists the general details of the alert on the General tab. For forwarding exceptions, additional details about the traffic flows affected by blackhole or traffic loss are displayed on the Metadata tab.

[Table on page 337](#) lists the fields on the Traffic Loss Alerts table.

Table 96: Fields on the Traffic Loss Alerts table

Field	Description
Severity	Severity of the alert: • Critical (Red hexagon icon)—Indicates blackhole. • Major (Orange square icon)—Indicates traffic loss or packet drops.
Start Time	The time when the packet drops or blackhole started occurring, in the MM-DD-YYYY HH:MM:SS format. For example, Jun 06, 2025 10:29:41 PM.

Table 96: Fields on the Traffic Loss Alerts table *(Continued)*

Field	Description
End Time	The time when the packet drop or blackhole was rectified, in the MM-DD-YYYY HH:MM:SS format. For example, Jun 06, 2025 10:29:41 PM.
Alert Name	Name of the alert raised by the device. For example, Packet Drops Observed, Blackhole Detected.
Exception	Message indicating the issue corresponding to the alert raised by the device. For example, Interface Block Incoming Drop.
Exception Type	Type of exception: • Interface • Host Interface • Fabric • Queuing • Forwarding
PFE	For forwarding exception types, the packet forwarding engine on the device that raised the alert.

By default, information for the past 30 minutes is displayed. To change this period, click the **Week**, **Day**, **3 hrs**, **1 hr**, **30 mins**, or **Custom** buttons provided above the graph.

Predictor Events

Predictor events indicate the routing, forwarding, and OS exceptions that are identified by Routing Director as a potential indicator of traffic loss.

To view predictor events on a device, enable the **Blackhole Detection** option present in the **Analytics** tab of the Device Profile assigned to the device in a network Implementation plan. See ["Add a Device Profile" on page 208](#) for details.

On the Predictor Events page, by default, information for the past 30 minutes is displayed. To change this period, click the **Week**, **Day**, **3 hrs**, **1 hr**, **30 mins**, or **Custom** buttons provided above the Predictor Events table.



NOTE: A maximum of 30 days of predictor events are displayed on the Predictor Events page.

Figure 33 on page 339 shows the Predictor Event page.

Figure 33: Predictor Events Page

The screenshot shows the 'Predictor Events' page with a table titled 'Predictor Events List (Potential Packet Loss)'. The table has columns for Severity, Start Time, End Time, Exception Type, Exception Code, Description, and PTE. There are four rows of data, all with a severity of 'Error'. The first row shows a packet loss event on May 27, 2025. The second and third rows show network path exceptions on May 23, 2025. The fourth row shows a packet loss event on May 23, 2025.

Severity	Start Time	End Time	Exception Type	Exception Code	Description	PTE
Error	May 27, 2025 00:00:00	May 27, 2025 00:00:01	Something	100	Pkt packet loss dropped	PPC0000
Error	May 23, 2025 10:00:00	May 23, 2025 10:00:01	Network	10000_PTE_02_0000	Network Path QoS	---
Error	May 23, 2025 10:00:00	May 23, 2025 10:00:01	Network	10000_PTE_02_0000	Network Path QoS	---
Error	May 23, 2025 10:00:00	May 23, 2025 10:00:01	Something	100	Pkt packet loss dropped	PPC0000

Table 97 on page 340 lists the fields on the Predictor Events table.

Table 97: Fields on the Predictor Events Table

Field	Description
Severity	Severity of the alert: • Critical—Indicates that an issue may be affecting the traffic flow and that needs immediate resolution. • Major—Indicates that an issue may be affecting the traffic flow that needs resolution. • Minor—Indicates that a non-critical issue may be affecting traffic flow. The event related to the alert should be monitored. • Warning—Indicates issues that may affect traffic flow if necessary actions are not taken. • Info—Indicates transient issues that may be potentially problematic.
Start Time	The start time (referenced to the local time zone) of the predictor event, in the MM-DD-YYYY HH:MM:SS format. For example, Jun 06, 2025 10:29:41 PM.
End Time	The end time (referenced to the local time zone) of the predictor event, in the MM-DD-YYYY HH:MM:SS format. For example, Jun 06, 2025 10:50:20 PM.
Exception Type	Type of exception: • Routing—Indicates that a route or next-hop is impacted which may lead to potential traffic loss. • Forwarding—Indicates packet flows are affected. • OS (Kernel)—Indicates host-path queues are impacted due to which traffic loss may occur.
Exception Code	Code for the exception.
Description	Description of the predictor event alert.

Table 97: Fields on the Predictor Events Table *(Continued)*

Field	Description
PFE	PFE on which the forwarding exception occurred. This field is not relevant for routing and OS exceptions.

For further information on a predictor event, select an alert and do any of the following:

- Click **View Details** to view details of the predictor event on the Predictor Event Details page.

The General tab of the Predictor Event Details page displays the general information (exception code, component on which the exception occurred, type of exception, and so on) and the Metadata tab displays flow information for forwarding exceptions and additional information for routing and OS exceptions.

- Click **View Correlated Events** to view events correlated to the predictor event. Correlated events identify other network events that have occurred on onboarded devices and that may have lead to the selected exception on the device.

[Figure 34 on page 341](#) shows the correlated events page.

Figure 34: Correlated Events



Table 98: Fields on the Correlated Events Table

Field	Description
Event ID	Identifier of the correlated event. A correlated event can be: • System log. To view details of the event related to the system log on the Events page (Observability > Health > Events), click the system log <i>event ID</i>. The Events page is filtered by the Event ID and you can view the details of the event related to the system log. • Exception (indicated as Ex in a balloon on the top-left corner of the event ID) For details of the exception, see Exception Code Reference. • Notification (indicated as NTF in a balloon on the top-left corner of the event ID)
Explanation	Explanation of how the event is correlated to the predictor event.
Confidence	Degree of confidence (relevance) that the correlated event has caused the predictor event. Confidence can be—Low, Medium, and High. Correlated events help to identify potential root cause of problematic exceptions or events that may lead to traffic loss. For example, IPv4 packet drop due to a discarded route could be because of a change in the device configuration or change in the route state. Correlated events are listed in the decreasing order of confidence.
Devices	Devices on which the correlated events have occurred.

Clicking the *Event ID* of a correlated event opens the Events page filtered by the event ID. Enable **Historical Data** to view additional events that may have contributed to the predictor event over a period of time. The correlated events table lists up to 10 correlated events in the decreasing order of relevance.



NOTE: Routing Director runs a job every 15 minutes to identify and list correlated events. As a result, it may take up to 15 minutes for a correlated event to be listed for a predictor event. For instance, if a correlated event is identified in the 1st minute of a 15-

minute interval when the job runs, the correlated event will be listed after 15 minutes. Conversely, if a correlated event is identified in the 11th minute of the 15-minute interval, it will be listed after 4 minutes.

SEE ALSO

[Juniper Resiliency Interface](#)

Device Connectivity Data and Tests Results

SUMMARY

This section provides an overview of connectivity tests, test results, and configurations that an administrator must perform to enable the tests.

IN THIS SECTION

- [Connectivity Accordion | 344](#)
- [Connectivity Details Page | 347](#)
- [View Connectivity Test Results | 350](#)

When you onboard a device, Routing Director automatically triggers test agents installed on your devices that generate synthetic traffic to initiate a connectivity test. The test streams run from a device to neighboring devices, edge routers, Internet endpoints (such as DNS service, HTTP service, and web services), and to the external hosts on Google Cloud Platform (GCP), Microsoft Azure, and Amazon Web Services (AWS) clouds. Routing Director supports connectivity tests to Asia, Europe, and North American regions of the three cloud providers. The duration of a connectivity test is one minute.

For details on devices that support test agents, see [Supported Junos OS Releases, Devices, and Browsers](#).

Configurations to Trigger Connectivity Tests

To enable Routing Director to initiate test connections during device onboarding, you must configure the interface profile. You can then associate the interface profile with one or more devices that you include in the network implementation plan. Users with Super User and Network Admin roles must perform the following configurations to enable Routing Director to initiate connectivity tests.

1. Internet Connected—Enable **Internet Connected** in the interface profile. When you include this interface profile in a network implementation plan, Routing Director triggers connectivity tests from specific or all device ports.

If you assign the interface profile as the default profile, Routing Director triggers Internet Endpoint and Cloud Provider connectivity tests on all ports of all devices that you configure in the network implementation plan. See ["Add an Interface Profile" on page 222](#) for more information.

2. Active Assurance—Configure device labels, endpoint device URLs, and the cloud provider hosts to which test agents run connectivity tests on the Create Device Profile) page. See ["Add a Device Profile" on page 208](#) for more information.

Connectivity Accordion

To access the Connectivity accordion, navigate to the **Observability > Health > Troubleshoot Devices** page. Click a device name to access the *Device-Name* page. Click the **Connectivity (accordion)** in the Overview tab.

The connectivity accordion on the *Device-Name* page displays the health of connections from a device in your network to a remote device. The accordion displays the overall status of the device connections at the top-right corner. The status displays **Urgent Action Needed** if critical events occur within the last seven days or **Healthy** if no connection issues are detected. You can view connection-specific details when you expand the accordion. You have the flexibility to configure automated connectivity tests when you plan device onboarding or use the Retest button in the connectivity accordion to run connectivity tests after onboarding devices. You can run the tests on all connections (ports) of devices or select connections on which you want to run the test.

After the tests are complete, you can view the results of these tests as links in the Connectivity accordion on the *Device-Name* page. [Table 100 on page 349](#) describes the fields in the Connectivity accordion.

Click the health status links for a connection to view details about the faulty connections on the ["Connectivity Details Page" on page 347](#). On the Connectivity Details page, you can rerun tests for specific or all remote endpoints after you resolve the connectivity issues.

The following list explains terms associated with connectivity tests:

- **Metrics**—Metrics such as delay, delay variance, HTTP timeout, ping (packet) loss, and round trip time (RTT) enable Routing Director to collect quantitative measurements to evaluate the quality of a connection.
- **Protocols**—Protocols such as HTTP, ping, and DNS are used to measure the metrics in a connectivity test. Ping is used to test connectivity from a device in your network to neighboring devices, edge devices, and known hosts in the cloud provider's network. HTTP and DNS protocols are used to test connectivity to Internet endpoints such as DNS service, HTTP service, or other web services.

- **Types of remote endpoints**—Types of remote devices to which test agents check connectivity. Remote endpoints can be neighboring devices, edge devices, Internet endpoints (DNS servers or web servers), and devices (external hosts) in the cloud.
- **Connection**—Connections are unidirectional flows of synthetic traffic from a test agent installed on a device to a test agent on another device, from a test agent to a reflector (BGP peering), or from a test agent to an external host in a public cloud. A connectivity test to a remote endpoint, such as neighboring devices, include multiple connections. Depending on the remote endpoint, each connection uses a protocol (such as ping) to check for select metrics (such as RTT). If a single connection (unidirectional traffic flow) experiences issues, the test fails.
- **Test Result**—Test Results are shown as timeline graphs of multiple key performance indicators (KPIs)—such as error seconds, response time, and packet loss—that indicate the health of a connection type. The KPIs are calculated based on the metric data collected for delay, delay variance, ping packet loss, round trip time, HTTP/ping response time, and HTTP timeout.



NOTE: Connectivity tests are not supported on third-party devices. As a result, the connectivity test results for third-party devices are not available on the Overview Tab.

Table 99: View Connectivity Information

Connection	Description
Neighbors	Neighbors are routers that use dynamic routing protocols to discover each other in a network topology. Neighbors can use multicast messages or unicast messages depending on the network configuration. Displays the number of neighboring devices connected to the device and the health of their connection (healthy or unhealthy) to the device.
Edges	Edge devices are devices at the perimeter that connect your network to another network. An edge device can be peering devices in your local network, an Internet Gateway, a customer edge (CE) or a provider edge (PE) device, an area border router (ABR), or an autonomous system border router (ASBR). Displays the number of edge devices (routers) connected to the device and the health of their connection (healthy or unhealthy) to the device.

Table 99: View Connectivity Information (*Continued*)

Connection	Description
Internet Endpoints	Endpoints are URLs that locate a service that is hosted on a remote server. Examples of services are HTTP service, DNS service, or other web services that you want to access. Displays the number of Internet endpoints (servers) and the health of their connection (healthy or unhealthy) to the device.
Cloud Providers	If you enable connectivity tests to public cloud providers in a device profile, test agents initiate a connectivity test from the device to known hosts in a public cloud provider's network. Displays the number of regions to which connectivity tests are initiated for Amazon Web Services, Google Cloud Platform, and Microsoft Azure. View the status of the connectivity (healthy or unhealthy) from the cloud host to the device.
Relevant Events	You can access events of varying severity that are generated for the tests within the last seven days. Click Details to view the device name, test description, and start time and end time of the tests. Click View All Relevant Events to open the Events for <i>Device-Name</i> page that displays all events generated for different connections from the device. The Events for <i>Device-Name</i> page contains the following information: • Severity—Displays Informational events for tests that pass and Critical events for tests that fail. • Time Stamp—Displays the date and time when test agents initiate the test. • Type—Displays the event type as Active Assurance. • Description—Specifies the test protocol, remote endpoint, and result of the test.
Retest	To re-run connectivity tests, click Retest and select All Connections if you want Routing Director to re-run connectivity tests on all the connections.

Connectivity Details Page

To access the Connectivity accordion, navigate to the **Observability > Health > Troubleshoot Devices** page. Click a device name to access the *Device-Name* page. Click the **Connectivity (accordion)** in the Overview tab. Click any hyperlink in the accordion and you are directed to the Connectivity Details page.

The Connectivity Details page contains the following sections:

- **Relevant Events**—After completing connectivity tests, Routing Director generates Critical and Information events for connectivity tests and bad cables. An Informational event denotes that the test passed, and a critical event denotes that the test failed. Click **View all Relevant Events** to view events triggered for all tests to the device's connections.
- **Refresh**—Routing Director automatically refreshes the data every 10 minutes and displays the time for the upcoming round of connectivity data refresh. Alternatively, click the **Refresh** icon to refresh the connectivity data for the device connections.
- **Show Connections Between**—Displays the category of remote endpoint (such as edge devices, neighboring devices, and cloud providers), and the number of devices in each remote endpoint category.

Enable the toggle button corresponding to a connection type to view the health of the connection in the topology view.

- **Connections Between Devices**—Displays a topology view of all connections from a device. For a remote endpoint, the topology view shows a single line that represents all connections from a device to multiple remote devices. You can perform the following tasks that are related to connections:
 - **Access details of faulty connections on the topology map**—After the connectivity tests are run on the onboarded device, the topology view displays the count of faulty connections. Faulty connections appear as red icons on the lines that indicate the connections. You can hover your cursor over the count icon to obtain details of the faults for a connection type.
 - **Run connectivity tests**—To re-run connectivity tests, click **Retest** and select **All Connections** if you want Routing Director to re-run connectivity tests on all the connections. Alternatively, you can select a specific connection (**Neighbors**, **Edges**, **Internet Endpoints**, or **Cloud Providers**) to which you want to re-run the test from the device.

After the test is complete, the topology view is automatically updated. In addition, the Connections table below the topology view displays the updated information.

To view additional details (such as a detailed view of the logs raised for events, errors, protocols used, and so on) of a test, click the details icon that appears when you hover your hover over the time range of the test.

- **Connections**—Displays a table with details about the connectivity tests run on the device. [Table 100 on page 349](#) describes the fields you see in the Connections table.
 - To view the test results, click the connectivity status (ERROR, PASSED, or FAILED) on the Connections table. The Test Results for *Device-name* to *Device-name* page appears. The Test Results page shows the KPIs and metrics collected from the test connections. You can view the results in timeline graphs. See "[View Connectivity Test Results](#)" on [page 350](#) for more information.
 - To view details of a connection, select a connection in the table. Click **More > Detail**. The **Source to Remote End Point** pane displays test time range, source, remote endpoint, test protocol, test result, and number of logs on the Details tab.

On the log tab, view log details for the connection such as the start time, end time, log level, and log message.

Table 100: Fields in the Connections Table

Field	Description
Status	Displays the status of the connection: • Scheduled—Displays SCHEDULED when a test agent is scheduled to trigger a connectivity test. • Running—Displays RUNNING when a test agent runs a connectivity test. • Waiting—Displays WAITING when the test agent is not ready for performing a test. For example, a user triggers a test but when the test agent is not available or offline, Routing Director displays WAITING until the default maximum timeout duration of 300 seconds and then displays the ERROR status. • Error—Displays ERROR when the test agent does not trigger a test connection. For example, when an interface goes down or the test agent goes offline, Routing Director displays the ERROR status. Hover over the error status to see the cause of the error. • Failed—Displays FAILED when an HTTP or ping connectivity test fails. The FAILED status is caused by test metrics such as delay, delay variance, or packet loss exceeding the threshold for a connection. • Passed—Displays PASSED if all connections to a remote endpoint is healthy (no errors or failures). To see more details, click the PASSED or FAILED health status link. The Test Result for <i>device-name-1</i> to <i>device-name-2</i> page appears where you can check the connectivity test results in detail. See "View Connectivity Test Results" on page 350 for more information.
Test Time Range	Displays the date and time range when a test is executed. The date is displayed in the DD/MM/YYYY format and the time as Minutes:Seconds, with the time zone.
Source	Displays the name of the device from which the synthetic traffic is sent.
Source Interface	Displays the name of the interface on the source device from which the synthetic traffic is sent.

Table 100: Fields in the Connections Table *(Continued)*

Field	Description
Destination	Name of the cloud provider and the region you previously configured for the connectivity test.
Remote End Point	Displays the name or management IP address of the remote device to which Routing Director initiates a connectivity test, along with its management interface name. Example: <i>Device-name:10.1.1.1</i>
Logs	Displays the number of logs generated for the connection from the device.
Protocol	Displays the protocol used for the test connection initiated by the test agent, such as HTTP, DNS, or Ping.
Type	Displays the type of remote endpoint to which the test agent initiates the connectivity test. For example, Internet Endpoints DNS, Cloud Endpoints Reachability, Edge Reachability, or Neighbor reachability.

View Connectivity Test Results

The **Test Results for *Device-name-1* To *Device-name-2*** page displays the KPIs and metrics as a timeline graph. KPIs in tests include error seconds, response time, and (packet) loss (%).

You can access the Test Results for *Device-name-1* To *Device-name-2* when you click **PASSED** or **FAILED** status links on the Connectivity Details page. Expand each KPI to view the measurements that are displayed in the timeline graph. [Table 101 on page 352](#) displays the test metrics for the DNS protocol, [Table 102 on page 353](#) displays the test metrics for the HTTP protocol, and [Table 103 on page 353](#) displays the test metrics for the ping protocol.

You can view the following information about the Test Results page:

- **Error Seconds**—For connectivity tests, the test results are compiled every 1 second. Every second, Routing Director checks if an error occurred for a test metric in a connection. If an error occurred in a connection, the Error Seconds graph displays 100%, else 0% for the test result compilation time. The graph plots the error seconds every time test results are compiled for the default test duration of 60 seconds.

Figure 35: Error Seconds Graph



- **Response time**—Displays the maximum, average, and minimum response time for DNS, HTTP, and ping packets in milliseconds. Delay variance or jitter is the variance in the amount of time that is taken by different packets when they traverse from a sending device to a receiving device. The less delay variance you measure in your network, the less latency you experience. Less latency is desirable in voice-based applications such as teleconferencing.

Figure 36: Response Time Graph



- **Loss**—The Percent lost metric measures the percentage of pings that are lost out of the total number of pings that are sent to a remote device. If the percentage of pings lost exceeds 60, then Routing Director generates a neighboring ping test alert.

Figure 37: Loss Graph



Table 101: Test Metrics for DNS Connectivity

Test Metrics	DNS
Error seconds (ES)	Yes
ES timeout	Yes
ES lifetime	Yes
ES response	Yes
ES loss	No
ES delay variance (DV)	No
Response Time	
Maximum response time (milliseconds)	Yes
Average response time (milliseconds)	Yes
Minimum response time (milliseconds)	Yes
Delay variance (milliseconds)	No
Loss	
Percent lost	No

Table 102: Test Metrics for HTTP Connectivity

Test Metrics	HTTP Request
ES timeout	Yes
ES lifetime	No
ES response	Yes
ES loss	No
ES delay variance (DV)	No
Response Time	
Maximum response time (milliseconds)	Yes
Average response time (milliseconds)	Yes
Minimum response time (milliseconds)	Yes
Delay variance (milliseconds)	No
Loss	
Percent lost	No

Table 103: Test Metrics for Ping Connectivity

Test Metrics	Ping Request
ES timeout	Yes
ES lifetime	No
ES response	Yes
ES loss	Yes
ES delay variance (DV)	Yes

Table 103: Test Metrics for Ping Connectivity *(Continued)*

Test Metrics	Ping Request
Response Time	
Maximum response time (milliseconds)	Yes
Average response time (milliseconds)	Yes
Minimum response time (milliseconds)	Yes
Delay variance (milliseconds)	Yes
Loss	
Percent lost	Yes

Device Management

IN THIS CHAPTER


- [Device Management Workflow | 355](#)
- [Device Licenses Overview | 357](#)
- [About the Features Tab | 358](#)
- [About the Licenses Tab | 360](#)
- [Manage Device Licenses | 361](#)
- [About the Software Images Page | 363](#)
- [Upload a Software Image | 366](#)
- [Delete a Software Image | 371](#)
- [About the Configuration Backups Page | 372](#)
- [Configuration Templates Overview | 375](#)
- [About the Configuration Templates Page | 376](#)
- [Add a Configuration Template | 379](#)
- [Edit and Delete a Configuration Template | 388](#)
- [Preview a Configuration Template | 389](#)
- [Deploy a Configuration Template to a Device | 390](#)

Device Management Workflow

Device life-cycle management (LCM) allows users with the Network Admin or Super User role to manage the onboarded devices in their organization. As part of device management, you can monitor how the onboarded devices are functioning, make configuration changes, and perform other tasks for the optimal performance of the devices.

[Figure 38 on page 356](#) displays the different tasks you can perform as part of the device management workflow.

Figure 38: Device Management Tasks



Before you start managing a device, verify that the device was successfully onboarded. If the status of the device on the Inventory page (**Inventory > Devices > Network Inventory**) is *Connected*, the device is onboarded successfully.

To manage a device, you can perform the following operations:

- Monitor the device health (including trust score) and the device performance by drilling-in to the various accordions. You can take actions based on the alerts and alarms that are generated for a device. See ["View Results of Automated Device Tests" on page 268](#).
- Manage the software image of a device. To ensure that there are no security-related vulnerabilities, it is important that you install the latest supported software image or, if required, upgrade the software image. See ["About the Software Images Page" on page 363](#).
- Back up the device configuration and restore the configuration if required. Configuration backups are useful because you can restore a device configuration in case of faulty configuration updates. See ["About the Configuration Backups Page" on page 372](#).
- Create customized configuration templates and deploy the configuration templates to devices. See ["About the Configuration Templates Page" on page 376](#).
- On the Troubleshoot Devices page (**Observability > Health > Troubleshoot Devices**), you can:
 - View device-related information (chassis, interfaces, licenses, and features) and add licenses for a device.
 - Assign the device to a site.
 - Reboot the device.

- Backup device configuration.
- Upgrade device software.
- Export device details in CSV format.

See ["About the Troubleshoot Devices Page" on page 454](#).



TIP: Routing Director provides the flexibility to apply new configurations and upgrade the software images for devices in two ways:

- Using configuration templates and software images that are available on Configuration Templates (**Inventory > Devices > Configuration Templates**) and Software Images pages (**Inventory > Devices > Software Images**) respectively. Use this approach if you want to apply new configurations to or update the software image for one device at a time.

For more information, see ["Add a Configuration Template" on page 379](#) and the *Upgrade the Image on a Device* section in ["About the Troubleshoot Devices Page" on page 454](#).

- Using an existing network implementation plan. You can edit an existing network implementation plan to update the device configurations and software images, and then publish the network implementation plan to apply the changes on devices that are part of the network implementation plan. Use this approach if you want to apply new configurations to or update the software images for more than one device at a time.

For more information, see ["Edit a Network Implementation Plan" on page 260](#) and ["Publish a Network Implementation Plan" on page 258](#).

RELATED DOCUMENTATION

[Device Life-Cycle Management Overview | 168](#)

[Device Onboarding Overview | 171](#)

Device Licenses Overview

You can add licenses to a device by using Routing Director, to use the features on the device. You need device licenses for each device in your network. Each device license is tied to a software feature. A

feature is a logical group of functionalities for a device that is specified with every license. For more information on licenses for ACX Series devices, see [Juniper Licensing User Guide](#).

After the device is onboarded, a superuser or network administrator can add a device license from the **Licenses** tab of the Routing Director GUI.

To add a license for a device, navigate to the **Observability > Health > Troubleshoot Devices > *device-name* > Inventory > Licenses** tab.

To view features of the device licenses that you already added for a device, navigate to the **Observability > Health > Troubleshoot Devices > *device-name* > Inventory > Features** tab.

You can do the following after you add a device license:

- View audit logs to confirm that the device licenses are correctly applied.
- Add more device licenses and remove the device licenses on the Licenses tab.
- View the device license inventory on the Licenses tab.
- View the licensed features per device on the Features tab.
- Filter the device license and feature information.

RELATED DOCUMENTATION

[About the Licenses Tab | 360](#)

[About the Features Tab | 358](#)

[Manage Device Licenses | 361](#)

About the Features Tab

IN THIS SECTION

- [Tasks You Can Perform | 359](#)
- [Field Descriptions | 359](#)

To access this page from the Routing Director GUI, click **Observability > Health > Troubleshoot Devices > *device-name* > Inventory > Features**.

A feature is a logical group of functionalities for a device. You must add a device license to use a feature on a device. You can use the Features tab to view the inventory of the features associated with the licenses added to the device.

Tasks You Can Perform

- View features related to the licenses added to the device. See [Table 104 on page 359](#).
- Sort, resize, or re-arrange columns in a table (grid).
- Filter the data displayed in the table—Click the filter icon (funnel) and select whether you want to show or hide advanced filters. You can then add or remove filter criteria, save criteria as a filter, apply or clear filters, and so on. The filtered results are displayed on the same page.

Field Descriptions

[Table 104 on page 359](#) describes the fields in the Features tabbed page:

Table 104: View Features of Device Licenses

Attribute	Description
Name	The name of the licensed feature.
Description	A short description of the licensed feature.
Used Count	The Number of licenses currently being used on the device. If a device license is added and the feature is configured, then the license is considered used.
Installed Count	The number of licenses installed on the device for the particular feature.
Need Count	The number of times a feature is used without a license.
End Date	Expiration information (date, time, and time zone) for the license. For example: 2012-03-30 01:00:00 IST.

About the Licenses Tab

IN THIS SECTION

- [Tasks You Can Perform | 360](#)
- [Field Descriptions | 360](#)

To access this page from the Routing Director GUI, click **Observability > Health > Troubleshoot Devices > *device-name* > Inventory > Licenses**.

Use this page to view details about the device licenses applied, and information about the number of features available for this device. You can add one or more than one device license for a device.

Tasks You Can Perform

- View the details of added device licenses. See [Table 105 on page 360](#) .
- Add a device license. See ["Add a Device License" on page 361](#).
- Delete a device license. See ["Delete a Device License" on page 362](#).
- Sort, resize, or re-arrange columns in a table (grid).
- Filter the data displayed in the table—Click the filter icon (funnel) and select whether you want to show or hide advanced filters. You can then add or remove filter criteria, save criteria as a filter, apply or clear filters, and so on. The filtered results are displayed on the same page.

Field Descriptions

[Table 105 on page 360](#) describes the fields in the Licenses tabbed page:

Table 105: View Details of Device Licenses

Attribute	Description
Name	The name (license key) of the device license.
Version	The numeric version number of the device license.

Table 105: View Details of Device Licenses (*Continued*)

Attribute	Description
State	The state of the license key. The following are the license key states: - Valid—The license key is valid. - Invalid—The license key that you have entered is invalid. For example, a license key can be invalid if you have entered an incorrect license key or if the license has expired. If a device license is invalid, alerts are generated to remind you to add a valid device license.
Software Serial Number	The software serial number of the device license. The software serial number is shipped electronically after you purchase a device license.
Features	The total number of features available with the device license. Click the link (total number displayed) to view the list of features.

Manage Device Licenses

SUMMARY

Read these topics to learn how you (Super User or Network Admin) can add and delete device licenses.

IN THIS SECTION

- [Add a Device License | 361](#)
- [Delete a Device License | 362](#)

Add a Device License

Add a device license to use a feature on a device.

To add a device license:

1. Navigate to **Observability > Health > Troubleshoot Devices > *device-name* > Inventory > Licenses**.

The Licenses tabbed page appears.

2. Click the add license (+) icon to add a new device license.

The Add License page appears.

3. Do one of the following:

- Click **Upload License** to upload the license file (.txt).

To upload the license file, click **Browse** and navigate to the license file (.txt) on your local file system, and click **Open**.



NOTE: Ensure that the license file is downloaded and saved in your local file system. You can download the license file from the Juniper Agile Licensing portal. You can also choose to receive the license file over an e-mail.

- Click **Enter License Details** and paste the license key that you copied or downloaded from the Juniper Agile Licensing portal.

4. Click **OK** to add the device license.

The device license is added. The Licenses tabbed page appears.

After you add a license, you can view the features included in the device license. Navigate to the **Observability > Health > Troubleshoot Devices > *device-name* > Inventory > Features** tab to view the inventory of features associated with the device license.

Delete a Device License

You can delete a device license by using the Routing Director GUI.



CAUTION: You can delete a device license even when you are using the licensed feature. When you delete a device license that is already in use,

- the license state becomes Invalid.
- alerts are generated to remind you to add a valid device license.

To delete a device license:

1. Navigate to **Observability > Health > Troubleshoot Devices > *device-name* > Inventory > Licenses**.

The Licenses tabbed page appears.

2. Select the device license that you want to delete and click the **Delete** (trash can) icon.

The device license is deleted. The Licenses tabbed page appears.

About the Software Images Page

IN THIS SECTION

- [Tasks You Can Perform](#) | 363
- [Field Descriptions](#) | 364

To access this page, click **Inventory > Devices Software Images**.

A software image is a software installation package that you use to upgrade or downgrade the OS running on a device. A superuser or a network administrator can upload the software images to Routing Director by using the GUI and upgrade or downgrade the images installed on the device.

Tasks You Can Perform

You can perform the following tasks from this page:

- View details of a software image—To view details of a software image, select the image name and click **More > Details**.

The Image Details pane appears on the right side of the page. The pane displays two sections:

- General—Displays the general information about the selected software image as listed in [Table 106 on page 364](#).
- Checksums—Displays all the checksums associated with the selected software image. The Checksums section includes values obtained from MD5, SHA1, SHA256, and SHA512 algorithms.
- View the basic information as listed in [Table 106 on page 364](#).

Click the **Close (x)** icon to close the Details pane.

- Upload a software image; see ["Upload a Software Image" on page 366](#).
- Delete a software image; see ["Delete a Software Image" on page 371](#).
- You can also perform the following tasks on this page:
 - Sort, resize, or re-arrange columns in a table (grid).
 - Show or hide columns in the table or reset page preferences, using the vertical ellipsis menu.

- Search by using keywords—Click the search icon (magnifying glass), enter the search term in the text box, and press Enter. The search results are displayed on the same page.
- Filter the data displayed in the table—Click the filter icon (funnel) and select whether you want to show or hide advanced filters. You can then add or remove filter criteria, save criteria as a filter, apply or clear filters, and so on. The filtered results are displayed on the same page.

For more information, see ["GUI Overview" on page 6](#).



NOTE: Currently, the filter option is not supported in Size, Description, Checksum, and Release Notes columns.

Field Descriptions

[Table 106 on page 364](#) displays the fields on the Software Images page.

Table 106: Fields on the Software Images Page

Field	Description
Image Name	Name of the software image file. For example, junos-evo-install-acx-f-x86-64-22.1R3.12-EVO.iso.
Version	Version number of the OS. Version number indicates the major software update that is released every quarter of a year. For example, Junos OS Evolved 23.1.
Release	Release number of the OS. Release number indicates the minor software update that addresses bugs and performance issues within a version. For example, Junos OS Evolved 22.3R2.
Vendor	Vendor of the device. Routing Director supports devices from Juniper Networks and Cisco Systems. See Supported Junos OS Releases, Devices, and Browsers for a complete list of devices supported by Routing Director.

Table 106: Fields on the Software Images Page *(Continued)*

Field	Description
Device series	Device model. For example, ACX7509, ACX7100. For the list of devices that Routing Director supports, see Supported Junos OS Releases, Devices, and Browsers.
Created By	User who uploaded the software image file.
Last modified	Date and time when the software image file was uploaded. The timestamp is displayed in the following format: Month DD, YYYY, HH:MM:SS AM/PM TIME ZONE. For example, June 14, 2023, 4:29:52 PM IST.
Size	Size of the software image.
Managed by	Type of user who manages the image file. The software images available in an organization are uploaded by either superusers or network administrators. For example: 'Org' when uploaded by a Super User or Network Admin of an organization.
Description	Description of the software image.
Checksum	Expected SHA256 checksum to validate the uploaded data. The images and their SHA256 checksum can be obtained from the Juniper Software Download page.
Release Notes	Click the link to view the release notes for the software image that you uploaded.

RELATED DOCUMENTATION

| [Device Life-Cycle Management Overview](#) | 168

Upload a Software Image

You must be a user with the Super User or a Network Admin role to upload a software image to Routing Director.

Before you upload a software image, ensure that you have downloaded the required software image from the [Juniper Software Download](#) page to your local system.

To upload a software image:

1. Click **Inventory > Devices > Software Images** in the left navigation menu.
The Software Images page appears.
2. Click the **Add (+)** icon.
The Upload Image page appears.
3. To upload the image from your local system, click **Browse** provided in the **Image File** field and select the image file that is saved on your computer in the explorer that opens.
To upload files larger than 2.5 GB, use the Upload Image REST API, instead of the GUI, See "[Sample Script](#)" on page 368 for a sample of the script that you can use to upload a software image by using the Upload Image REST API.



NOTE: If you get a message indicating that an error has occurred while uploading an image, see "[Increase VM Disk Space for Software Image Upload](#)" on page 370 for resolution.

4. Enter values by referring to [Table 107 on page 366](#).
5. Click **OK**.

The image is uploaded to the Routing Director database and listed on the Software Images page.

A message indicating that the upload is successful is displayed with a link to the Audit Logs page (**Systems Menu > Audit Logs** present on the GUI banner). You can view the progress of the upload on the Audit Logs page.

Table 107: Fields on the Upload Images Page

Field	Description
Image File*	Click Browse to upload an image file to from your local system. Before you upload the image file, ensure that you have downloaded the required software image from the Juniper Software Downloads page, and you have saved the software image in your local system.

Table 107: Fields on the Upload Images Page *(Continued)*

Field	Description
Display Name*	Enter or modify the software image name. The Display Name field is automatically populated with the name of the file you uploaded. If the name of the software image file uploaded does not meet the naming criteria, you can modify or enter a name. The name can contain only alphanumeric characters, hyphen, period, underscore, and the maximum length allowed is 254 characters.
Description	Enter the description for the software image file.
Vendor*	Select the vendor of the device from the drop-down list. Routing Director supports devices from Juniper Networks and Cisco Systems.
Device Series*	Select the model of the device from the drop-down list.
Release*	Enter the release number of the OS. Release number indicates the minor software update that addresses bugs and performance issues within a version. For example, Junos OS Evolved 22.3R1.
Version*	Enter the version number of the OS. Version number indicates the major software update that is released every quarter of a year. For example, Junos OS 23.1.
Expected SHA256*	Provide the expected SHA256 checksum to validate the uploaded data. The images and their SHA256 checksum can be obtained from the Juniper Software Download page. On the Software Download page, select the device from the Find a Product box. Click and expand the Install Package section and select the software package and checksums for the release.

Table 107: Fields on the Upload Images Page (*Continued*)

Field	Description
Release Notes Link	Provide the link to the release notes. You can find the link to the release notes on the Juniper Software Download page. On the Software Download page, select the device from the Find a Product box. Click and expand the Documentation and Release Notes section and find the release notes link for the release.

Sample Script to Upload a Software Image to Routing Director

The following is a sample script to upload a software image to Routing Director. Enter appropriate values for:

- *host*—The hostname of the server from where the software image is to be uploaded to Routing Director.
- *user*—The username to access the server.
- *password*—The password to access the server.
- *image-filename*—The name of the software image to be uploaded.
- *path*—The path where the software image is stored on the server.

```
import base64
import hashlib
import json
import requests
URL = 'https://<host>/api/v1/devicesoftware/'
USER = '<user>'
PASSWORD = '<password>'
ORG_ID = 'ff118da5-000a-483c-91a1-45d478478548'
MIN_SIZE = 5 * 1024 * 1024
FILENAME = '<image_filename>'
DIRECTORY = '<path>'
def file_sha(contents):
    h = hashlib.sha256()
    data = contents.read(MIN_SIZE)
    while data:
        h.update(data)
        data = contents.read(MIN_SIZE)
    return h.hexdigest()
```

```

def print_sw_images():
    response = requests.get(URL + ORG_ID + '/images', auth=(USER, PASSWORD), verify=False)
    images = decode(response)
    print('There are %d images' % len(images))
    for img in images:
        print(img.get('name'))
def decode(response):
    if response.status_code != 200:
        print(response.status_code)
        print(response.text)
        exit(1)
    return json.loads(response.text)
def upload(file_name: str, data, file_sha256: str, vendor: str, series: str, release: str,
version: str):
    headers = {'Content-Type': 'application/json'}
    url = URL + ORG_ID + '/upload'
    response = requests.post(url, auth=(USER, PASSWORD), headers=headers, verify=False, json={
        'name': file_name,
        'expected_sha256': file_sha256,
        'device_vendor': vendor,
        'device_series_list': [series],
        'release': release,
        'version': version,
    })
    r = decode(response)
    upload_id = r.get('id')
    print('Uploading %s' % upload_id)
    url = url + '/' + upload_id
    file_slice = read_min(data)
    while file_slice:
        encoded = base64.b64encode(file_slice).decode('utf-8')
        response = requests.put(url, auth=(USER, PASSWORD), headers=headers, verify=False, json={
            'file_data': 'data:application/octet-stream;base64,' + encoded
        })
        msg = decode(response)
        print(msg)
        file_slice = read_min(data)
    response = requests.put(url, auth=(USER, PASSWORD), headers=headers, verify=False, json={
        "complete": True
    })
    msg = decode(response)
    print(msg)
def read_min(file):

```

```

ret = b''
while len(ret) < MIN_SIZE:
    file_slice = file.read(MIN_SIZE)
    ret = ret + file_slice
    if not file_slice:
        break
return ret
if __name__ == '__main__':
    print_sw_images()
    with open(DIRECTORY + FILENAME, 'rb') as iso:
        sha = file_sha(iso)
        print(sha)
    with open(DIRECTORY + FILENAME, 'rb') as iso:
        upload(FILENAME, iso, sha, 'Juniper', 'ACX7100', '22.2R1.12-EV0', '22.2')
    print_sw_images()

```

This requirements file can be used to install any dependencies necessary-

```
requests==2.28.1
```

Increase VM Disk Space for Software Image Upload

The default quota for image upload is limited to 40% of the total Ceph disk size divided by 3. For example, if you are using minimum recommended 50GiB Ceph disk space per VM, then the quota for image upload is around 27GiB.

To check the current quota for image upload, use:

```

root@controller:~# kubectl get cephobjectstore -n rook-ceph object-store -o
jsonpath='{.spec.dataPool.quotas}' | jq
{
  "maxSize": "27Gi"
}

```

To check the current usage of the quota, use:

```

root@controller:~# paragon-ceph-usage-check
Checking S3 bucket usage
Listing top-level directories and checking usage:
S3 Directory list : ['devicesoftware', 'usw2-jcloud-dev-paa-plugin-service-plugins-storage']
Usage for: devicesoftware
Total Objects: 0

```

```
Total Size: 0 Bytes
Usage for: usw2-jcloud-dev-paa-plugin-service-plugins-storage
Total Objects: 0
Total Size: 0 Bytes
```

If you get a message stating that an error has occurred while uploading an image, increase the maximum size of the Ceph object Storage daemon (OSD) pool to a higher value (for example 30GiB) by executing the following command in the installer node of the Routing Director cluster:

```
root@controller:~# kubectl patch cephobjectstore -n rook-ceph object-store --type=merge -p
'{"spec": {"dataPool": {"quotas": {"maxSize": "30Gi"}}}}'
cephobjectstore.ceph.rook.io/object-store patched
```



NOTE: It may take a few minutes for the changes to be applied.

To confirm that the Ceph OSD pool size is increased to 30GiB, execute the following command:

```
root@controller:~# kubectl exec -ti -n rook-ceph $(kubectl get po -n rook-ceph -l app=rook-ceph-
tools -o jsonpath={..metadata.name}) -- ceph osd pool get-quota object-store.rgw.buckets.data
quotas for pool 'object-store.rgw.buckets.data':
max objects: N/A
max bytes : 30 GiB (current num bytes: 32212254720 bytes)
```

The `max bytes` indicates the Ceph OSD pool size in GiB and bytes.

You can use the software uploaded to Routing Director to upgrade the image installed on devices. See ["Upgrade Software" on page 308](#) or ["Upgrade Device Image" on page 458](#).

RELATED DOCUMENTATION

[Device Life-Cycle Management Overview](#) | 168

Delete a Software Image

You can delete one or more software images from the Software Images page when you no longer need them.

To delete a software image, you must be assigned either a Super User or and Network Admin role.

To delete a software image:

1. Click **Inventory > Devices > Software Images**.

The Software Images page appears.

2. Select one or more software images that you want to delete and click the **Delete** icon (trashcan).

A confirmation message appears.

3. Click **Yes** to delete the software images.

The selected images are deleted and the images are no longer listed on the software Images page.

RELATED DOCUMENTATION

[Device Life-Cycle Management Overview | 168](#)

[About the Software Images Page | 363](#)

About the Configuration Backups Page

IN THIS SECTION

● [Tasks You Can Perform | 372](#)

● [Field Descriptions | 374](#)

To access this page, click **Inventory > Devices > Configuration Backups**.

In Routing Director, a superuser or network administrator can take a back up of configuration from a device and restore the configuration on the device, if required. Faulty configuration updates can cause emergencies leading to network outages. If you have taken a backup, you can rollback to the back up in the event of a faulty configuration update. This rollback helps you to resume normal operations.

The Configuration Backup page lists the configuration files that are generated when you back up a device from the Troubleshoot Devices (**Observability > Troubleshoot Devices**) page.

Tasks You Can Perform

You can perform the following tasks from this page:

- View the list of device configuration files that were backed-up and their details.

To view the details of a specific configuration file, select a device configuration file. Click **More > Details** or hover over the device configuration file and click the **Detailed View** icon. The Configuration Backup Details pane appears on the right side of the page displaying the device configuration details. For more information, see [Table 108 on page 375](#).

Click the **Close (x)** icon to close the pane.



NOTE: You can view the list of configuration files only if you have backed up a device from the Troubleshoot Devices (**Observability > Troubleshoot Devices**) page.

- Preview device configuration.

You can preview the device configuration file before you push the configuration to a device.

To preview a device configuration, select a device configuration file from the list and click the **Preview** button. A Preview Configuration page appears displaying the device configuration in both JSON and XML formats.



TIP: On the Configuration accordion (**Observability > Troubleshoot Devices > Device-Name**), you can compare an active version of the configuration committed on a device against other backed-up versions of the same device. For more information, see ["Configuration Data and Test Results" on page 308](#).

- Restore a device configuration.

If you have a device configuration backup, you can restore the device configuration in case of faulty device configuration updates. The restore operation enables you to restore a previously-saved device configuration.

To restore a device configuration:

1. Select the device configuration from the list and click the **Restore Configuration** button.

The Select Device to Restore page appears with a list of devices. The device from which the backup was taken is selected by default. The list is filtered based on the device family of the backed-up configuration you have selected.



WARNING:

- Restore a device configuration only if you have selected the correct device from which the configuration backup was taken.
- Before you restore a device, you must verify the device name and model from the device family list against the device name and model of the selected backed-up configuration.
- The restore operation fails if you select an incorrect device to restore a configuration.

2. Select the device for which you want to restore the device configuration and click **OK**.

A confirmation message appears stating that the backup configuration is successfully restored.

- You can also perform the following tasks on this page:
 - Sort, resize, or re-arrange columns in a table (grid).
 - Show or hide columns in the table or reset page preferences, using the vertical ellipsis menu.
 - Search by using keywords—Click the search icon (magnifying glass), enter the search term in the text box, and press Enter. The search results are displayed on the same page.
 - Filter the data displayed in the table—Click the filter icon (funnel) and select whether you want to show or hide advanced filters. You can then add or remove filter criteria, save criteria as a filter, apply or clear filters, and so on. The filtered results are displayed on the same page.

For more information, see ["GUI Overview" on page 6](#).

- Set the number of rows to be displayed in the table—Set the number of rows (15, 30, 50, or 100) to be displayed in the Network Implementation Plan table in the **Display** drop-down list provided at the bottom right-corner of the table.

If there are multiple pages of rows, use the arrow heads to traverse through the pages or enter a page number to view the entries listed on that page.

- Refresh the Configuration Backups table and set auto refresh frequency—A Refresh icon is provided at the bottom left-corner of the Configuration Backups table. Click the **Refresh** icon to view the most recent updates made to the network implementation plans.

Click the **ellipsis** next to the Refresh icon to set the auto refresh frequency. You can either turn off the auto refresh of the table or set the auto refresh frequency to 30 seconds, 1 min, 3 mins, and 5 mins.

Field Descriptions

[Table 108 on page 375](#) displays the fields on the Configuration Backups page.

Table 108: Fields on the Configuration Backups Page

Field	Description
Name	The name of the configuration backup file.
Created	The date and time when the configuration backup was taken. The timestamp is displayed in the following format: Month DD, YYYY, HH:MM:SS AM/PM TIME ZONE. For example, May 8, 2023, 4:29:52 PM IST.
Device	The hostname of the device from which the backup was taken.
Model	The model of the device. For example, ACX7024. For the list of devices that Routing Director supports, see Supported Junos OS Releases, Devices, and Browsers.
OS	The OS installed on the device. For example, Junos or Junos Evolved.
Operator	The e-mail address of the user who backed up the configuration.

RELATED DOCUMENTATION

[Configuration Data and Test Results](#) | 308

[About the Troubleshoot Devices Page](#) | 454

Configuration Templates Overview

IN THIS SECTION

- [Benefits](#) | 376

In Routing Director, you can use configuration templates to provision customized configurations throughout the device life-cycle for Juniper Networks devices.

Using configuration template, superusers and network administrators can create customized configuration, preview the configuration template, and deploy the configuration to one or more devices. You can view, add, edit, or delete configuration templates from the Configuration Templates (**Inventory > Devices > Configuration Templates**) page.

You can apply a configuration template to all devices or to specific device in a network.

Benefits

- Using configuration templates, you can create customized configurations and push the configurations to one or more devices. This helps you to deploy additional configurations beyond the standard configuration templates provided in Routing Director.
- As your network grows, you can deploy existing configuration templates to new devices effortlessly. You can push the same configurations to one or more devices, reducing the chances of manual errors and inconsistencies.
- Configuration templates enable standardization and adherence to best practices, leading to enhanced network security and reliability.

RELATED DOCUMENTATION

| [About the Configuration Templates Page](#) | 376

About the Configuration Templates Page

IN THIS SECTION

- [Tasks You Can Perform](#) | 377
- [Field Descriptions](#) | 378

To access this page, click **Inventory > Devices > Configuration Templates**.

Using configuration templates, you can create customized configuration and deploy the configuration to one or more devices. For more information, see "[Configuration Templates Overview](#)" on page 375.

**NOTE:**

Tasks You Can Perform

You can perform the following tasks from this page:

- View details of a configuration template.

Select a configuration template and click **More > Details** or hover over the configuration template and click the **Detailed View** icon. The Details of *Configuration-Template-Name* pane appears. The pane displays two tabs:

- **GENERAL**— Displays the general information about the selected configuration template as listed in [Table 109 on page 378](#).
- **TEMPLATE**— Displays the template you defined in CLI or XML format in the Template Configuration tab (**Inventory > Devices > Configuration Templates > Add (+)**).
- Preview a configuration template; see ["Preview a Configuration Template" on page 389](#).
- Deploy a configuration template on one or more devices; see ["Deploy a Configuration Template to a Device" on page 390](#).
- Add a configuration template; see ["Add a Configuration Template" on page 379](#).
- Edit and delete a configuration template; see ["Edit and Delete a Configuration Template" on page 388](#).
- You can also perform the following tasks on this page:
 - Sort, resize, or re-arrange columns in a table (grid).
 - Show or hide columns in the table or reset page preferences, using the vertical ellipsis menu.
 - Search by using keywords—Click the search icon (magnifying glass), enter the search term in the text box, and press Enter. The search results are displayed on the same page.
 - Filter the data displayed in the table—Click the filter icon (funnel) and select whether you want to show or hide advanced filters. You can then add or remove filter criteria, save criteria as a filter, apply or clear filters, and so on. The filtered results are displayed on the same page.

For more information, see ["GUI Overview" on page 6](#).

Field Descriptions

Table 109 on page 378 displays the fields on the Configuration Templates page.

Table 109: Fields on the Configuration Templates Page

Field	Description
ID	The ID of the configuration template. An ID is assigned to a configuration template at the time of its creation.
Name	The name of the configuration template.
Type	The format in which the configuration template is defined—CLI, NETCONF EDIT, NETCONF RPC, or Non EXECUTABLE.
Family	The device family for which the configuration template is applicable: • ACX • EX • JNP • JUNIPER ANY • MX • NFX • PTX • QFX
Description	A description of the configuration template.
Last Updated	The date and time when the configuration template was last updated, in the Month DD, YYYY, HH:MM:SS AM/PM TIME ZONE format. For example, May 8, 2023, 4:29:52 PM IST.

Table 109: Fields on the Configuration Templates Page (*Continued*)

Field	Description
Created by	The user who created the configuration template. If the column displays <i>System</i>, it indicates that the configuration template is a predefined configuration template.

Add a Configuration Template

You must be a user assigned either the Super User or the Network Administrator role with the privilege to add configuration templates.

To add a configuration template:

1. Select **Inventory > Devices > Configuration Templates**.

The Configuration Templates page appears.

2. Click the **Add** icon (+).

The Add Configuration Template page appears.



NOTE: Fields marked with an asterisk (*) are mandatory.

3. Configure the fields on the Basic Information tab according to the guidelines provided in [Table 110 on page 382](#).
4. Click **Next** to go to the Template Configuration tab.
5. Add the configuration on the Template Configuration tab.

You can do the following in the editor provided for entering the configuration:

- Copy the required configuration stanza from a device and create a template from parameters in the configuration.
- Parameterize variables by using double curly braces `{{}}`.
- You can view a sample configuration by clicking the **Sample Configuration** link. The sample that you see depends on the Device Vendor (Juniper or Cisco) that you selected on the Basic Information Tab.

The following are the sample configurations that are available for Juniper and Cisco devices:

- Sample configuration to configure interfaces on a Juniper device by using CLI:

```
configure
{% if interfaces and interfaces | length > 0 %}
{% for interface in interfaces %}
    set interfaces {{interface.name}} vlan-tagging
    {% for ifl in interface.ifls %}
        set interfaces {{interface.name}} unit {{ifl.unit}} vlan-id {{ifl.vlan_id}}
        {% for family in ifl.families %}
            set interfaces {{interface.name}} unit {{ifl.unit}} family {{family.name}}
            address {{family.address}}
        {% endfor %}
    {% endfor %}
{% endfor %}
{% endif %}
commit
```

- Sample configuration to configure interfaces on a Juniper device by using NETCONF EDIT and NETCONF RPC:

```
<edit-config>
<target><candidate /></target>
<default-operation>merge</default-operation>
<config>
  <configuration>
    <interfaces>
      {% for interface in interfaces %}
        <interface>
          <name>{{ interface.name }}</name>
          <vlan-tagging/>
          {% for ifl in interface.ifls %}
            <unit>
              <name>{{ifl.unit}}</name>
              <vlan-id>{{ifl.vlan_id}}</vlan-id>
              <family>
                {% for family in ifl.families %}
                  <{{family.name}}>
                    <address>
                      <name>{{family.address}}</name>
                    </address>
                  </{{family.name}}>
                {% endfor %}
              </family>
            </unit>
          {% endfor %}
        </interface>
      {% endfor %}
    </interfaces>
  </configuration>
</config>
```

```

    {% endfor %}
  </family>
</unit>
  {% endfor %}
</interface>
{% endfor %}
</interfaces>
</configuration>
</config>
</edit-config>

```

- Sample configuration to configure timezone on a Cisco device.

```

{
  "replace": [
    {
      "path": {
        "elem": [
          {
            "name": "openconfig-system:system"
          },
          {
            "name": "clock"
          }
        ]
      },
      "val": {
        "json_ietf_val": "{{ietf_value}}"
      }
    }
  ]
}

```



NOTE: In the NETCONF EDIT and NETCONF RPC configurations, Routing Director assigns values only for interface name. For other parameters such as VLAN ID and address family, you must provide suitable values in the configuration. Routing Director will not automatically assign any values.

For Service Orchestration, configure interface name in the `{{ interface.name.split("")[0] }}` format. For example:

```
<interface>
  <name>{{ interface.name.split(".")[0] }}</name>
  <vlan-tagging/>
  <unit>
    <name>{{ interface.name.split(".")[1] }}</name>
    <vlan-id>1</vlan-id>
    <family>
      <inet>
        <address>
          <name>10.10.10.10/24</name>
        </address>
      </inet>
    </family>
  </unit>
</interface>
```

6. Click **Next** to go to the Generated UI tab. You can view the parameters you set in the Template Configuration tab in the Generated UI tab.
7. Perform one or more actions on the Generated UI tab, as explained in [Table 111 on page 383](#).
8. Click **Save**.

A message indicating that the configuration template is added is displayed with a link to the Audits Logs page (**Settings Menu > Audit Logs**).



NOTE: Ensure that the size of the configuration template file is not more than 1 MB to ensure smooth operation of device onboarding and service provisioning.

[Table 110 on page 382](#) lists fields to be entered on the Basic Information tab of the Add Configuration Templates page.

Table 110: Fields on the Basic Information Tab of the Add Configuration Templates Page

Field	Description
Template Name*	Enter a unique name for the configuration template. The name can only contain alphanumeric characters, hyphen, period, and underscore; 64-characters maximum.
Description	Enter a description for the configuration template; 255-characters maximum.

Table 110: Fields on the Basic Information Tab of the Add Configuration Templates Page (Continued)

Field	Description
Device Vendor*	Select a device vendor from the drop-down list: • Juniper (default) • Cisco
Configuration Format*	Select the output format for the configuration template. The below options are listed if you select a Juniper device: • CLI (default) • NETCONF EDIT • NETCONF RPC • NON EXECUTABLE OpenConfig/Cisco UM mixed is the only configuration format supported for a Cisco device.
Device Family*	Select a device family for which you are deploying the configuration template: • Juniper device family includes ACX, EX, JNP, JUNIPER ANY, MX, NFX, PTX, QFX, and SRX. • Cisco device family includes 8000, ASR, Cisco-ANY, XRV, and NCS.

[Table 111 on page 383](#) lists the actions that you can perform on the Generated UI tab of the Add Configuration Templates page.

Table 111: Generated UI Actions (Add Configuration Template Page)

Action	Description
Reorder the UI	Drag and drop individual fields, grids, or sections to change the order in which the parameters appear on the UI.

Table 111: Generated UI Actions (Add Configuration Template Page) (Continued)

Action	Description
Modify the settings for a field, section, or grid	To modify the settings for a field, section, or grid: 1. Click the Settings (gear) icon next to the field, section, or grid. The Form Settings pane appears on the right side of the page, displaying the Basic Settings and Advanced Settings tabs. 2. Modify the fields on these tabs, as needed. See Table 112 on page 385 for an explanation of the fields on these tabs. 3. Click Save Settings for each field to save your changes. The modifications that you made are displayed on the UI.
Reset the generated UI	Click Undo all Edits to discard the changes that you made and undo the changes made on the UI.
Preview configuration	Preview the configuration defined in the configuration template. To preview a configuration template: 1. Click Preview Configuration. The Preview Configuration page appears, displaying the configuration that was rendered based on the values that you entered. 2. Check if the configuration was rendered correctly. • If the configuration was not rendered correctly, click the close (X) icon to go back to the Generated UI tab and make modifications as needed. • If the configuration was rendered correctly, click OK. You are returned to the Generated UI page.

[Table 112 on page 385](#) lists the fields on the Form Settings pane.

Table 112: Form Settings (Add Configuration Template Page)

Setting	Guideline
<i>Basic Settings Tab</i>	Fields populated in this tab are based on the input type that you select for a parameter.
Input Type	The input type for the parameter in the configuration template. Select: • Text (default): If the input value for the parameter is a string of characters. • Number: If the input value for the parameter is a number. • Email: If the input value for the parameter is an e-mail address. • IPv4: If the input value for the parameter is an IPv4 address. • IPv4 Prefix: If the input value for the parameter is an IPv4 prefix. • IPv6: If the input value for the parameter is an IPv6 address. • IPv6 Prefix: If the input value for the parameter is an IPv6 prefix. • Toggle Button for Boolean: If the input value for the parameter is a Boolean value (true or false). The field is enabled if the • Dropdown: If the input value for the parameter is selected from a list. • Password: If the input value for the parameter is a password. The value that you enter is masked (default). (Optional) Click the Show Password (eye) icon to unmask the password. • Confirm Password: If the input value for the parameter is to confirm the password. If you select this option, a Confirm Password field appears on the UI. The value that you enter is masked (default). (Optional) Click the Show Password (eye) icon to unmask the password.
Label	Enter the label that you want displayed (on the UI) for the parameter.
Default Value	Specify a default value for the parameter.

Table 112: Form Settings (Add Configuration Template Page) (Continued)

Setting	Guideline
Validate	For Text input type, select one or more validation criteria against which the input value will be checked: • No Space • Alpha and Numeric • Alpha, Numeric, and Dash • Alpha, Numeric, and Underscore If the value that you entered for the parameter on the UI does not meet the selected validation criteria, an error message appears. NOTE: For greater control of input values, you can use the regular expression option in the Advanced Settings tab.
Description	Enter an explanation for the parameter, which will appear when you hover over the Help (?) icon for the parameter; the maximum length allowed is 256 characters.
Global Scope	Enable the toggle button to make the parameter common across all devices to which the configuration template is being deployed. A Global (G) icon is displayed beside the selected parameter. If you disable the toggle button, which is default, the parameter must be specified for each device.
Hidden	Click the toggle button to hide the parameter on the UI when you preview and deploy the template. Typically, this option is used to hide a parameter and display it in the template only when an event is triggered. By default, the toggle button is disabled, which means that the parameter is displayed.
Required	Click the toggle button to make the parameter mandatory; parameters that are mandatory are marked with an asterisk (*) on the UI.
Maximum Value	For parameters that are numbers, enter the maximum value (up to 16 digits) for the input.

Table 112: Form Settings (Add Configuration Template Page) (Continued)

Setting	Guideline
Minimum Value	For parameters that are numbers, enter the minimum value (up to 16 digits) for the input.
Visibility for Disabled	For Boolean parameters, select one or more parameters that must appear on the UI when the toggle button is disabled (Boolean value is FALSE).
Visibility for Enabled	For Boolean parameters, select one or more parameters that must appear on the UI when the toggle button is enabled (Boolean value is TRUE).

Table 112: Form Settings (Add Configuration Template Page) (Continued)*Advanced Settings Tab*

Regex	Enter a regular expression (regex pattern) to validate the input value. A regular expression defines a search pattern that is used to match characters in a string. For example, the regular expression [A-Z] matches the input with the characters A through Z. If the input consists of characters other than A through Z, an error message (as specified in the Invalid Message field) appears.
Invalid Message	Enter an error message that you want to display on the UI when the input value does not match the specified regular expression.

What's Next

You can deploy the template on devices; see ["Deploy a Configuration Template to a Device" on page 390](#).

RELATED DOCUMENTATION

[Device Life-Cycle Management Overview | 168](#)

[About the Troubleshoot Devices Page | 454](#)

[Device Licenses Overview | 357](#)

Edit and Delete a Configuration Template

IN THIS SECTION

- [Edit a Configuration Template | 388](#)
- [Delete a Configuration Template | 388](#)

You must be a user that is assigned the Super User or Network Admin role with privileges to edit and delete configuration templates.

Edit a Configuration Template

To edit a Configuration Template:

1. Select **Inventory > Devices > Configuration Templates**.

The Configuration Templates page appears.

2. Select the configuration template that you want to modify and click the **Edit** (pencil) icon.

The Edit Configuration Template page appears. The fields on this page are the same as the fields that you configure in the Add Configuration Template workflow.

3. Modify the fields as needed.

Refer ["Add a Configuration Template" on page 379](#) for an explanation of the fields.



NOTE: Fields marked with an asterisk (*) are mandatory.

4. Click **OK**.

The modifications are saved and you are returned to the Configuration Templates page, where a confirmation message is displayed. If the configuration template was previously deployed on a device, then you must redeploy the configuration template for the changes to take effect.

Delete a Configuration Template



NOTE: You can delete a configuration template only if the following conditions hold good:

- You added (created) the template.

- The template is not deployed on a device.

1. Select **Inventory > Devices > Configuration Templates**.

The Configuration Templates page appears.

2. Select the configuration template that you want to delete and click the **Delete** (trash can) icon.

You are asked to confirm the delete operation.

3. Click **Yes**.

You are returned to the Configuration Templates page, and a pop-up appears indicating whether the deletion was successful or not.

SEE ALSO

[Device Life-Cycle Management Overview | 168](#)

[Device Licenses Overview | 357](#)

[About the Troubleshoot Devices Page | 454](#)

Preview a Configuration Template

You must be a user that is assigned the Super User or Network Admin role with the preview privilege to preview configuration templates.

You can use the Preview option to validate a configuration template. You can enter values for the parameters in the configuration template and then render the template to view the configuration.

To preview a configuration template:

1. Select **Inventory > Devices > Configuration Templates**.

The Configuration Templates page appears.

2. Select the configuration template that you want to preview and click **Preview**.

The Template Preview for *Configuration-Template-Name* page appears.

3. In the CONFIGURE tab, specify values for the parameters as needed.



NOTE: Fields marked with an asterisk (*) are mandatory.

4. After you have entered the necessary parameters, click **PREVIEW**.

The PREVIEW tab generates the configuration based on the values that you specified.

5. Check if the configuration was rendered correctly.

If the configuration was not rendered correctly, you can modify the configuration template as needed. See ["Edit and Delete a Configuration Template" on page 388](#).

6. Click **Close**.

You are returned to the Configuration Templates page. You can deploy configuration templates on a device.



TIP: You can preview a configuration template only if the configuration you entered in the Template Configuration tab is correct and complete. If you preview when the template configuration is not correct or complete, a yellow alert icon is displayed adjacent to the name of the selected template. The alert displays a warning indicating that the configuration template is incomplete and asks you to edit the selected configuration template.

RELATED DOCUMENTATION

[Device Life-Cycle Management Overview | 168](#)

[About the Troubleshoot Devices Page | 454](#)

Deploy a Configuration Template to a Device

You can deploy a configuration template on one or more devices in an organization. This operation enables you to apply new configurations to devices after a device is onboarded or to deploy additional configurations to a device.

You must be a user assigned either the Super User or Network Admin role to be able to deploy configuration templates on devices.

To deploy a configuration template to one or more devices:

1. Select **Inventory > Devices > Configuration Templates**.

The Configuration Templates page appears.

2. Select the configuration template that you want to deploy and click **Deploy to Devices**.

A Deploy *template name* to Devices page appears listing the devices to which the configuration template can be assigned.

3. Select one or more devices and click **Deploy**.

The Set Configuration Template Parameters page appears.

- a. In the Configure tab, assign values for the parameters.

- b. Click **Preview** to view and generate the configuration.

If the configuration is fine, click **OK** or change the configuration in the Preview tab if you want to change the configuration.

4. Click **Deploy**.

The settings that you entered are saved and you are returned to the Configuration Templates page. A message indicating that the deployment is successful is displayed with a link to the Audit Logs page (**Administration > Audit Logs**). You can view the progress of the deploy operation in the Audit Logs page.

RELATED DOCUMENTATION

[Device Life-Cycle Management Overview | 168](#)

[About the Troubleshoot Devices Page | 454](#)

4

PART

Observability

- [Introduction | 393](#)
 - [Health Dashboard | 398](#)
 - [Troubleshoot Devices | 450](#)
 - [Monitor Devices Using AI/ML | 488](#)
 - [View Network Topology | 498](#)
 - [Routing Observability | 565](#)
 - [Rules | 576](#)
-

Introduction

IN THIS CHAPTER

- [Observability Overview | 393](#)

Observability Overview

IN THIS SECTION

- [Network Observability and Topology Visualization | 394](#)
- [Routing Observability | 394](#)
- [Device Observability and Troubleshooting | 395](#)
- [Health Dashboard | 395](#)
- [Network Events Observability | 396](#)
- [Benefits | 396](#)

The observability use case in Routing Director enables you to view your entire network topology, monitor network health, and get notifications of anomalies in the network and network devices. Observability enables you to get actionable insights into the health of your network to identify and to remediate network issues.

Routing Director uses telemetry obtained from the network, network devices, and network services to monitor and understand what is happening across the network. Routing Director then detects and helps resolve real-time issues to keep the network and its components healthy, delivering an easily sustainable high-quality experience for network operators. Observability is a critical use case of Routing Director.

With observability, Network Operations Center (NOC) engineers (typically superusers, network administrators, and observers) have fine-grained visibility and ability to view the network topology, and monitor the health and quality of the network and its components (such as devices, links, and routing protocols) through a single pane of glass UI. Additionally, drill-in views provide detailed information

about all network components. Routing Director monitors and analyzes network components by using key performance indicators (KPIs), logs, and metrics, and notifies you about network issues through alerts and alarms. You can also choose to receive notifications of network issues using webhooks and over e-mail, which enables you to continue to monitor the network even when you are not logged in to the GUI.

You can use observability for effective day 2 operations of the network. Day 2 operations focus on observing the state of the network and its components and guiding the actions that a network administrator performs on the basis of alerts, alarms, and device system logs to maintain a healthy and operational network.

We explain the different categories of the observability use case in this topic.

Network Observability and Topology Visualization

Routing Director provides NOC engineers an up-to-date view of the Layer 3 IP network. You can observe the network and visualize the topology by using the topology map feature. The topology map in the GUI provides an intuitive and multidimensional view of the network topology including sites, devices, tunnels, and links. In addition, you can also customize the topology map for better visibility and analysis of the network topology. For more information, see ["Network Visualization Options" on page 501](#).

The network information table displays detailed information about network elements in the topology. You can view information about all active devices and the sites where they are deployed, details about the links between the devices, the alerts generated on the devices and sites, and the tunnels in your network. You can drill-down to individual devices and troubleshoot the alert conditions to maintain a healthy network. For more information, see ["Network Table Overview" on page 515](#).

In addition, Routing Director preconfigures and installs test agents on the devices that it manages. Test agents generate synthetic traffic that verify connectivity between the devices in your network topology. You can analyze the results of these connectivity tests to determine and fix any issues.

Routing Observability

Routing Director allows you to actively monitor the overall routing health in real time by using the Route Explorer dashboard. You can use the Route Explorer dashboard to monitor the health and status of all BGP routes and prefixes. You can use this information to identify routers that are down, prefix-related changes, and troubleshoot connectivity issues. For more information, see ["About the Route Explorer Page" on page 566](#).

Device Observability and Troubleshooting

Routing Director provides you a detailed view of the devices and the connectivity between the devices in your network. You can monitor the health of the devices in the network and view all anomalies (that may need user intervention) in the devices.

Routing Director runs a series of automated health checks on the network and the network devices. The results of these checks provide a granular view of the health of hardware, software, interfaces, routing, compliance with Center for Internet Security (CIS) benchmarks, and connectivity of a device. You can drill down to a device to view information, such as:

- Remote management data
- CPU and memory utilization, fans, and PSUs
- Available physical interfaces, nonoperational interfaces, and data on input and output traffic
- Information on the SIRT advisories and genuineness of the OS installed on the devices
- Location, version, and compliance of the committed configuration with CIS benchmarks
- Device connectivity health and data
- Routing protocols, and health information related to BGP, OSPF, IS-IS, RSVP, LSP, and LDP neighbors

For information on drill-in views, see ["About the Device-Name Page" on page 461](#).

Routing Director detects and generates alerts and alarms for issues in your devices and in your overall network. You can also get notifications for alerts and alarms over external applications such as e-mail and Slack. Timely detection of issues enables you to fix them immediately and minimize the impact of such issues on your network and its performance. You can monitor these alerts and alarms from a single page, and drill down to the devices, to easily identify and fix the issues generating them. For more information, see ["About the Troubleshoot Devices Page" on page 454](#).

You can view a list of all alert and alarm events and choose to monitor specific events using event templates. In addition and importantly, you can subscribe to be notified of specific events over e-mail and Slack.

You can analyze device system logs to further analyze the status and health of the devices in the network. For more information, see ["About the Events Page" on page 470](#).

Health Dashboard

Routing Director provides a health dashboard, which displays a comprehensive view of metrics. The metrics help you to understand the overall health of your network. The Health Dashboard has the following tabs:

- WAN Health—Provides data on the health of the devices, interfaces, and routing components in your network. For more information, see ["About the WAN Health Tab" on page 399](#).
- Trust—Provides the average trust score of the targets, and data on device compliance and vulnerabilities in your network. For more information, see ["About the Trust Tab" on page 416](#).
- Active Assurance—Provides data on the overall health of measurements in your network. For more information, see ["About the Active Assurance Tab" on page 434](#).

You can use these metrics to understand the behavior of your network, analyze the impact of potential issues, and take data-driven actions to fix the issues.

Network Events Observability

Routing Director monitors Key Performance Indicators (KPIs) for a device, connectivity between the devices, and your overall network, and generates alerts when any anomalies are detected. Alerts are generated for interface, hardware, routing, and connectivity issues in your network. You can view these alert events and subscribe to be notified of the events over e-mail and over Slack. For more information, see ["Alerts Tab" on page 471](#).

Benefits

- Identify performance degradation of the network using the results of health-checks, alerts, and alarms
- Monitor the overall health of your network using Health Dashboard, which presents real-time data that is easy to understand
- Monitor routing data in real time to help identify routing issues and take corrective actions
- Get a centralized view of your network topology to help in network planning
- Detect device faults and network issues and send alert and alarm notifications over e-mail and Slack for quick resolution of issues
- Boost network operational efficiency and deliver high-quality experience for network operators
- Maintain low operational costs

RELATED DOCUMENTATION

[Routing Observability Overview | 565](#)

[Troubleshoot Using Alerts and Alarms | 450](#)

[Network Implementation Plan Overview | 230](#)

Health Dashboard

IN THIS CHAPTER

- [About the Health Dashboard | 398](#)
- [About the WAN Health Tab | 399](#)
- [About the Devices Health Page | 402](#)
- [About the Interfaces Health Page | 406](#)
- [About the Routing Health Page | 411](#)
- [About the Trust Tab | 416](#)
- [About the Vulnerability Health Page | 419](#)
- [About the Compliance Health Page | 423](#)
- [About the Integrity Health Page | 427](#)
- [Update Trust KPIs in an Air-Gapped Installation | 430](#)
- [About the Active Assurance Tab | 434](#)
- [About the Applications Page | 435](#)
- [About the L3 Network Page | 439](#)
- [About the L2 Network Page | 442](#)
- [About the TV & Multicast Page | 446](#)

About the Health Dashboard

IN THIS SECTION

- [Tasks You Can Perform | 399](#)

To access the Health Dashboard page, click **Observability > Health > Health Dashboard**.

Use the Health Dashboard page to monitor network health-related data. You can view and monitor WAN health, network trust information, and the health of various Active Assurance measurements in your network in real time.

Tasks You Can Perform

You can perform the following tasks from this page:

- View WAN Health-related information—Click the WAN Health tab to view information about the overall health of devices, interfaces, and routing neighbors in your network in real time. See [About the WAN Health Tab](#).
- View Trust-related information—Click the Trust tab to view the trust score, the trust plan applied, overall trust trend, compliance of devices, and information about vulnerabilities that affect overall device health. See [About the Trust Tab](#).
- View Active Assurance-related information—Click the Active Assurance tab to monitor the overall health of various measurements in your network in real time. See [About the Active Assurance Tab](#).
- Expand or collapse accordions.
 - Click **Expand All** to expand the accordions on a tab. You can view detailed information related to KPIs that affect overall health based on alerts, and a performance graph.
 - Click **Collapse All** to collapse all accordions on a tab. You can only view overall health that is based on alerts, and a performance graph.
- Update the dashboard or the tab(s) on the dashboard by clicking the refresh icon.

About the WAN Health Tab

IN THIS SECTION

- [Task You Can Perform | 400](#)

To access this tab, click **Observability > Health > Health Dashboard > WAN Health**.

This tab provides a unified view of the overall health of devices, interfaces, and routing neighbors in your network in real-time. You can monitor overall health and use this health-related data to perform corrective actions when necessary.

The WAN Health tab consists of the following accordions:

- Devices Health
- Interfaces Health
- Routing Health

Each accordion has three sections. The first section provides a description of the accordion. The second section displays the latest aggregate health snapshot. The last section displays a health trend graph representing the average health for the past 30 minutes.

Task You Can Perform

You can perform the following tasks from this page:

- View device health information on the Devices Health accordion.

On the Devices Health accordion, you can:

- View a health snapshot of the average health (in percentage) of devices, and the total number of unhealthy devices. The data displayed for healthy devices is calculated based on the number of devices that have an alert raised divided by the installed base. Install base indicates the total number of devices that are available in the network.
 - View the percentage and health of the KPIs, and the total number of unhealthy devices in relation to that KPI that affect device health.
 - Click any KPI (listed under Feature) or total number of devices (listed under Unhealthy Devices) to view more information about the performance of that KPI on the Devices Health page. See ["About the Devices Health Page" on page 402](#).
 - View a health trend graph of the average health trend of the devices for the past 30 minutes. Hover over any point in the graph to view the average health score for that category at that time.
- View interfaces health information on the Interfaces Health Accordion.

On the Interfaces Health accordion, you can:

- View a health snapshot of the average health (in percentage) of the interfaces of the devices, and the total number of unhealthy devices.
- View the percentage and health of the KPIs that affect interface health, and the total number of unhealthy devices in relation to that KPI that affect interface health.

- Click any KPI (listed under Feature) or total number of devices (listed under Unhealthy Devices) to view more information about the performance of that KPI on the Interfaces Health page. See ["About the Interfaces Health Page" on page 406](#).
- View a health trend graph of the average health trend of the interfaces for the past 30 minutes. Hover over any point in the graph to view the average health score for that category at that time.
- View routing health information on the Routing Health Accordion.

On the Routing Health accordion, you can:

- View a health snapshot of the average routing health (in percentage) in relation to the number of devices, and the total number of unhealthy devices.
- View the percentage and health of the KPIs, and total number of unhealthy devices in relation to that KPI that affect routing health.
- Click any KPI (listed under Feature) or total number of devices (listed under Unhealthy Devices) to view more information about the performance of that KPI on the Routing Health page. See ["About the Routing Health Page" on page 411](#).
- View a health trend graph of the average routing health trend for the past 30 minutes. Hover over any point in the graph to view the average health score for that category at that time.
- Expand or collapse accordions.
 - Click **Expand All** to expand the accordions on the tab. You can view detailed information related to KPIs that affect overall health based on alerts, and a performance graph.
 - Click **Collapse All** to collapse all accordions on the tab. You can only view overall health that is based on alerts, and a performance graph.
- View the total number of devices in the network in the top-right corner of the tab.
- View the timestamp of the last update or refresh (in *DD/MM/YYYY hh:mm:ss AM/PM* format) in the top-right corner of the tab. For Example: 07/12/2024 12:42:43 PM.
- Update the details on the tab by clicking the refresh icon.

RELATED DOCUMENTATION

| [About the Health Dashboard](#) | 398

About the Devices Health Page

IN THIS SECTION

- [Tasks You Can Perform | 402](#)
- [Sections on the Page | 403](#)

You can access this page in the following ways:

- Click **Observability > Health > Health Dashboard > WAN Health (Tab) > Devices Health (Accordion) > View Details**.

You can view detailed information about the overall health of devices in your network.

- Click **Observability > Health > Health Dashboard > WAN Health (Tab) > Devices Health (Accordion) > Feature (click any KPI listed)**.

You can view information about the overall device health in relation to the KPI that you selected.

The Devices Health page consists of a gauge chart, list of KPIs, Overview tab, Timeline tab, Affected Items tab, and Distribution tab.

Tasks You Can Perform

You can perform the following tasks from this page:

- View a gauge chart of the overall device health.
- View a list of four KPIs that affect device health. The KPIs listed are Temperature, Power, CPU, and Memory. You can select more KPIs from the **Select Others** drop down list available at the end of KPIs list.
- On the Overview tab:
 - View the percentage of healthy and unhealthy devices in the donut chart. The total number of devices in the network is displayed in the middle of the donut chart.
 - View the percentage of healthy and unhealthy devices grouped by OS versions in a stacked bar chart.
 - View the percentage of healthy and unhealthy devices grouped by device models in a single stacked bar chart.

- On the Timeline tab:
 - View the percentage of healthy devices for a predefined or a custom period on a graph. The graph is updated based on the KPIs that you select.
- On the Affected Items tab:
 - View the details of unhealthy devices in the order of severity of the alerts generated on those devices.
 - Click any device hostname to view alerts for that device on the Alerts for *Device-Hostname* page.
- The Distribution tab consists of the OS tab, Device Model tab, and Sites tab.
 - On the OS tab, view the details of unhealthy devices grouped by device OS version.
 - On the Device Model tab, view the details of unhealthy devices grouped by device model.
 - On the Sites tab, view the details of unhealthy devices grouped by sites.
- Click **Clear Selection** to clear all selected KPIs.
- View the total number of devices in the network on the top-right corner of the page.
- View the time stamp of the last update or refresh (in *MMM DD, YYYY hh:mm:ss AM/PM* format) in the top-right corner of the tab. For Example: May 12, 2024 12:42:43 PM
- Update the page by clicking the refresh icon.
- On the Affected Items and Distribution tab, you can also:
 - Sort, resize, or re-arrange columns in a table (grid).
 - Show or hide columns in the table or reset page preferences, using the vertical ellipsis menu.
 - Search by using keywords—Click the search icon (magnifying glass), enter the search term in the text box, and press Enter. The search results are displayed on the same page.
 - Filter the data displayed in the table—Click the filter icon (funnel) and select whether you want to show or hide advanced filters. You can then add or remove filter criteria, save criteria as a filter, apply or clear filters, and so on. The filtered results are displayed on the same page.

Sections on the Page

The following topics describe the sections on the Devices Health page.

Gauge Chart

The gauge chart displays the overall health of network devices as a percentage. Hover over the chart to view the percentage of healthy devices in relation to device health. The percentage of healthy devices is also displayed below the chart. The percentage is 0 even if one device in the network is unhealthy or has an alert associated to it.

List of KPIs

KPIs such as Temperature, Power, CPU, and Memory are listed next to the gauge chart. The overall health of a KPI is displayed next to the KPI itself. Click any KPI to view more factors that affect the overall health of that KPI.

For example, click **Temperature** to view other KPIs (such as PSM temperature, Routing Engine temperature, Chassis temperature, and FPC temperature) that contribute to the overall percentage that is displayed. In some instances, you can also view a second level of KPIs that affect the overall temperature percentage. For example, click **Routing Engine** to view CPU temperature, which is another related KPI that affects the overall temperature.

A total of four KPIs are listed by default next to the gauge chart. At the end of the KPI list, you can select more KPIs from the **Select Others** drop-down list to view additional KPIs such as State, RPM, FPC, and NTP. You can click any KPI to view more factors that affect the overall health. To clear all selected KPIs, click **Clear Selection**.

Overview Tab

The Overview tab consists of the Devices, OS Versions, and Models widgets.

The Devices widget displays the overall percentage of healthy and unhealthy devices in a donut chart. However, when you select a KPI from the list next to the gauge chart, the Devices widget displays the overall health of the device in relation to that KPI. The total number of devices in the network is displayed in the middle of the chart. Healthy devices are highlighted in green in the donut chart, and unhealthy devices are highlighted in red. Hover over the green portion of the donut chart to view the percentage and total number of healthy devices. Hover over the red portion of the donut chart to view the percentage and total number of unhealthy devices.

The OS Versions widget displays the percentage of healthy and unhealthy devices grouped by OS versions in a stacked bar chart. Hover over the bar chart to view this percentage. The portion of the bar chart highlighted in green represents healthy devices. The portion of the bar chart highlighted in red represents unhealthy devices. A maximum of ten OS versions are displayed in the widget in the descending order of the unhealthy device percentage.

The Models widget displays the percentage of healthy and unhealthy devices grouped by device models in a stacked bar chart. Hover over the bar chart to view this percentage. The portion of the bar chart

highlighted in green represents healthy devices. The portion of the bar chart highlighted in red represents unhealthy devices. A maximum of ten device models are displayed in the widget in the descending order of the unhealthy device percentage.

Timeline Tab

The Timeline tab displays the percentage of healthy devices for a predefined (a week, a day, 3 hours, 1 hour, 30 minutes) or a custom period on a graph. By default, performance for the past 30 minutes is displayed. To change this period, click the **Week**, **Day**, **3 hrs**, **1 hr**, **30 mins**, or **Custom** buttons provided above the graph.

You can also use the range picker provided above the graph to view information for a predefined or custom period. You can click the arrows provided on either side of the range picker to move between ranges. For example, click the **Week** button to view device health information for a week. Click the arrow on the left of the range picker to view information for the previous week. Click the arrow on the right of range picker to view information for the following week.

Data lines for the overall device health and device health in relation to individual KPIs (Temperature, Power, CPU, Memory, and so on) are displayed on the graph by default. Hover over any point on the line graph to view the percentage health for that particular point. When you select a KPI from the list next to the gauge chart, the graph is updated to only show the overall health of the device in relation to that KPI. The KPIs are also provided below the graph with colors automatically assigned to each of the KPIs. Click any KPI to view a data line in the assigned color on the graph. Click **Clear Selection** to clear all selected KPIs.

Affected Items Tab

The Affected Items tab consists of the Devices tab.

The Devices tab displays details of unhealthy devices in the order of severity of the alerts generated on those devices. The various alerts are:

- **Action Needed**—The device's hardware and temperature have issues that you must address.
- **Urgent Action Needed**—The device's hardware and temperature have issues that must be addressed immediately.

Click any device hostname to view alerts for that device on the Alerts for *Device-Hostname* page.

Distribution Tab

The Distribution tab consists of the OS, Device Model, and Sites tabs.

The OS tab displays the details of unhealthy devices grouped by the device OS version. The details on this tab are updated based on the KPIs that you select. You can view the total number of affected

devices, total number devices that use the OS version, and the percentage impact for each OS version on the tab. Click the link in the **Devices Affected** column to view details of all affected devices that use that OS version, on the Devices Affected for *OS-Name* page. Click the link in the **Total Network** column to view details of all devices that use that OS version, on the Total Devices for *OS-Name* page.

The Device Model tab displays the details of unhealthy devices grouped by device (hardware) model. The details on this tab are updated based on the KPIs that you select. You can view the total number of affected devices, total number of devices that belong to the hardware model category, and the percentage impact for each hardware model category on the tab. Click the link in the **Devices Affected** column to view details of all affected devices that belong to that hardware model category, on the Devices Affected for *Model-Name* page. Click the link in the **Total Network** column to view all devices that belong to that hardware model category, on Total Devices for *Model-Name* page.

The Sites tab displays the details of unhealthy devices grouped by sites. A site is the location of a device in an organization. The details on this tab are updated based on the KPIs that you select. You can view the total number of affected devices, total number devices in the site, and the percentage impact for each site on the tab. Click the link in the **Devices Affected** column to view details of all affected devices that belong to that site, on the Devices Affected for *Site-Name* page. Click the link in the **Total Network** column to view details of all devices that belong to that site, on the Total Devices for *Site-Name* page.

RELATED DOCUMENTATION

[About the WAN Health Tab | 399](#)

About the Interfaces Health Page

IN THIS SECTION

- [Tasks You Can Perform | 407](#)
- [Sections on the Page | 408](#)

You can access this page in the following ways:

- Click **Observability > Health > Health Dashboard > WAN Health (Tab) > Interfaces Health (Accordion) > View Details**.

You can view detailed information about the overall health of devices (based on interface health) in your network.

- Click **Observability > Health > Health Dashboard > WAN Health (Tab) > Interfaces Health (Accordion) > Feature (click any KPI listed)**.

You can view information about the overall health of devices based on the interface-related KPIs that you selected.

The Interfaces Health page consists of a gauge chart, list of KPIs, Overview tab, Timeline tab, Affected Items tab, and Distribution tab.

Tasks You Can Perform

You can perform the following tasks from this page:

- View a gauge chart of the overall device health.
- View a list of four KPIs related to interfaces that affect device health. The KPIs listed are Link Failure, RX Errors, TX Errors, and Link Flaps. You can select more KPIs from the **Select Others** drop down list available at the end of KPIs list.
- On the Overview tab:
 - View the percentage of healthy and unhealthy devices in the donut chart. The total number of devices in the network is displayed in the middle of the donut chart.
 - View the percentage of healthy and unhealthy devices based on interface health grouped by OS versions in a single stacked bar chart.
 - View the percentage of healthy and unhealthy devices based on interface health grouped by device models in a single stacked bar chart.
- On the Timeline tab:
 - View the percentage of healthy devices (based on interface health) for a predefined or a custom period on a graph. The graph is updated based on the KPIs that you select.
- On the Affected Items tab:
 - View the details of unhealthy devices (based on interface health) in the order of severity of the alerts generated on those devices.
 - Click any device hostname to view alerts for that device on the Alerts for *Device-Hostname* page.
- The Distribution tab consists of the OS tab, Device Model tab, and Sites tab.
 - On the OS tab, view the details of unhealthy devices (based on interface health) grouped by device OS version.

- On the Device Model tab, view the details of unhealthy devices (based on interface health) grouped by device model.
- On the Sites tab, view the details of unhealthy devices (based on interface health) grouped by sites.
- Click **Clear Selection** to clear all selected KPIs.
- View the total number of devices in the network on the top-right corner of the page.
- View the time stamp of the last update or refresh (in *MMM DD, YYYY hh:mm:ss AM/PM* format) in the top-right corner of the tab. For Example: May 12, 2024 12:42:43 PM
- Update the page by clicking the refresh icon.
- On the Affected Items and Distribution tab, you can also:
 - Sort, resize, or re-arrange columns in a table (grid).
 - Show or hide columns in the table or reset page preferences, using the vertical ellipsis menu.
 - Search by using keywords—Click the search icon (magnifying glass), enter the search term in the text box, and press Enter. The search results are displayed on the same page.
 - Filter the data displayed in the table—Click the filter icon (funnel) and select whether you want to show or hide advanced filters. You can then add or remove filter criteria, save criteria as a filter, apply or clear filters, and so on. The filtered results are displayed on the same page.

Sections on the Page

The following topics describe the sections on the Interfaces Health page.

Gauge Chart

The gauge chart displays the overall health of network devices, as a percentage. Hover over the chart to view the percentage of healthy devices based on interface health. The percentage of healthy devices based on interfaces is also displayed below the chart. The percentage is 0 even if one interface in the network is unhealthy or has an alert associated to it.

List of KPIs

KPIs related to Link Failure, RX Errors, TX Errors, and Link Flaps that affect interface health are listed next to the gauge chart. The overall health of a KPI is displayed next to the KPI itself. Click any KPI to view more factors that affect the overall health.

For example, click **TX Errors** to view other KPIs (such as Input and CRC) that contribute to the overall percentage that is displayed. In some instances, you can also view a second level of KPIs that affect the overall percentage. For example, click **CRC** to view Output CRC, which is another related KPI that affects the overall health.

A total of four KPIs are listed by default next to the gauge chart. At the end of the KPI list, you can select more KPIs from the **Select Others** drop-down list to view additional KPIs such as TX, RX, Optical, Temperature, and Errors. You can click any KPI to view more factors that affect the overall health. To clear all selected KPIs, click **Clear Selection**.

Overview Tab

The Overview tab consists of the Devices, OS Versions, and Models widget.

The Devices widget displays the overall percentage of healthy and unhealthy devices (based on interface health) in a donut chart. However, when you select a KPI from the list next to the gauge chart, the Devices widget displays the overall health of the device in relation to that KPI. The total number of devices in the network is displayed in the middle of the chart. Healthy devices are highlighted in green in the donut chart, and unhealthy devices are highlighted in red. Hover over the green portion of the donut chart to view the percentage and the total number of healthy devices. Hover over the red portion of the donut chart to view the percentage and the total number of unhealthy devices.

The OS Versions widget displays the percentage of healthy and unhealthy devices (based on interface health) grouped by OS versions in a stacked bar chart. Hover over the bar chart to view this percentage. The portion of the bar chart highlighted in green represents healthy devices. The portion of the bar chart highlighted in red represents unhealthy devices. A maximum of ten OS versions are displayed in the widget in the descending order of the unhealthy device percentage.

The Models widget displays the percentage of healthy and unhealthy devices (based on interface health) grouped by device models in a stacked bar chart. Hover over the bar chart to view this percentage. The portion of the bar chart highlighted in green represents healthy devices. The portion of the bar chart highlighted in red represents unhealthy devices. A maximum of ten device models are displayed in the widget in the descending order of the unhealthy device percentage.

Timeline Tab

The Timeline tab displays the percentage of healthy devices (based on interface health) for a predefined (a week, a day, 3 hours, 1 hour, 30 minutes) or a custom period on a graph. By default, performance for the past 30 minutes is displayed. To change this period, click the **Week**, **Day**, **3 hrs**, **1 hr**, **30 mins**, or **Custom** buttons provided above the graph.

You can also use the range picker provided above the graph to view information for a predefined or custom period. You can click the arrows provided on either side of the range picker to move between ranges. For example, click the **Week** button to view interface health information for a week. Click the

arrow on the left of the range picker to view information for the previous week. Click the arrow on the right of range picker to view information for the following week.

Data lines for the overall interface health and interface health in relation to individual KPIs (Link Failure, RX Errors, TX Errors, Link Flaps, and so on) are displayed on the graph by default. Hover over any point on the line graph to view the percentage health for that particular point. When you select a KPI from the list next to the gauge chart, the graph is updated to only show the overall interface health in relation to that KPI. The KPIs are also provided below the graph with colors automatically assigned to each of the KPIs. Click any KPI to view a data line in the assigned color on the graph. Click **Clear Selection** to clear all selected KPIs.

Affected Items Tab

The Affected Items tab consists of the Devices tab.

The Devices tab displays details of unhealthy devices (based on interface health) in the order of severity of the alerts generated on those devices. The various alerts are:

- **Action Needed**—The interfaces have issues that you must address (may not be immediately).
- **Urgent Action Needed**—The interfaces have issues that must be addressed immediately.

Click any device hostname to view alerts for that device on the Alerts for *Device-Hostname* page.

Distribution Tab

The Distribution tab consists of the OS, Device Model, and Sites tabs.

The OS tab displays the details of unhealthy devices (based on interface health) grouped by the device OS version. The details on this tab are updated based on the KPIs that you select. You can view the total number of affected devices, total number devices that use the OS version, and the percentage impact for each OS version on the tab. Click the link in the **Devices Affected** column to view details of all affected devices that use that OS version, on the Devices Affected for *OS-Name* page. Click the link in the **Total Network** column to view details of all devices that use that OS version, on the Total Devices for *OS-Name* page.

The Device Model tab displays the details of unhealthy devices (based on interface health) grouped by device (hardware) model. The details on this tab are updated based on the KPIs that you select. You can view the total number of affected devices, total number of devices that belong to the hardware model category, and the percentage impact for each hardware model category on the tab. Click the link in the **Devices Affected** column to view details of all affected devices that belong to that hardware model category, on the Devices Affected for *Model-Name* page. Click the link in the **Total Network** column to view all devices that belong to that hardware model category, on Total Devices for *Model-Name* page.

The Sites tab displays the details of unhealthy devices (based on interface health) grouped by sites. A site is the location of a device in an organization. The details on this tab are updated based on the KPIs that you select. You can view the total number of affected devices, total number devices in the site, and the percentage impact for each site on the tab. Click the link in the **Devices Affected** column to view details of all affected devices that belong to that site, on the Devices Affected for *Site-Name* page. Click the link in the **Total Network** column to view details of all devices that belong to that site, on the Total Devices for *Site-Name* page.

RELATED DOCUMENTATION

[About the WAN Health Tab | 399](#)

About the Routing Health Page

IN THIS SECTION

- [Tasks You Can Perform | 412](#)
- [Sections on the Page | 413](#)

You can access this page in the following ways:

- Click **Observability > Health > Health Dashboard > WAN Health (Tab) > Routing Health (Accordion) > View Details**.

You can view detailed information about the overall health of devices (based on routing health) in your network.

- Click **Observability > Health > Health Dashboard > WAN Health (Tab) > Routing Health (Accordion) > Feature (click any KPI listed)**.

You can view information about the overall health of devices based on the routing-related KPIs that you selected.

The routing health is calculated based on the following routing components:

- The states of all BGP, OSPF, IS-IS, RSVP, LSP, and LDP neighbors.
- LSP state and flaps.

- Entries available in the routing information base (RIB), and forwarding information base (FIB).

The Routing Health page consists of a gauge chart, list of KPIs, Overview tab, Timeline tab, Affected Items tab, and Distribution tab.

Tasks You Can Perform

You can perform the following tasks from this page:

- View a gauge chart of the overall device health.
- View a list of four routing health-related KPIs that affect overall device health. The KPIs listed are BGP, IS-IS, OSPF, RIB, MPLS, and FIB. You can select more KPIs from the **Select Others** drop down list available at the end of the KPIs list.
- On the Overview tab:
 - View the percentage of healthy and unhealthy devices in the donut chart. The total number of devices in the network is displayed in the middle of the donut chart.
 - View the percentage of healthy and unhealthy devices in relation to routing health grouped by OS versions in a single stacked bar chart.
 - View the percentage of healthy and unhealthy devices in relation to routing health grouped by device models in a single stacked bar chart.
- On the Timeline tab:
 - View the percentage of healthy devices (based on the health of routing components such as BGP, IS-IS, OSPF, RSVP, LDP neighbors; LSP state and flaps; and the state of RIB and FIB) for a predefined or a custom period on a graph. The graph is updated based on the KPIs that you select.
- On the Affected Items tab:
 - View the details of unhealthy devices (based on routing health) in the order of severity of the alerts generated on those devices.
 - Click any device hostname to view alerts for that device on the Alerts for *Device-Hostname* page.
- The Distribution tab consists of the OS tab, Device Model tab, and Sites tab.
 - On the OS tab, view the details of unhealthy devices (based on routing health) grouped by device OS version.
 - On the Device Model tab, view the details of unhealthy devices (based on routing health) grouped by device model.

- On the Sites tab, view the details of unhealthy devices (based on routing health) grouped by sites.
- Click **Clear Selection** to clear all selected KPIs.
- View the total number of devices in the network on the top-right corner of the page.
- View the time stamp of the last update or refresh (in *MMM DD,YYYY hh:mm:ss AM/PM* format) in the top-right corner of the tab. For Example: May 12, 2024 12:42:43 PM
- Update the page by clicking the refresh icon.
- On the Affected Items and Distribution tab, you can also:
 - Sort, resize, or re-arrange columns in a table (grid).
 - Show or hide columns in the table or reset page preferences, using the vertical ellipsis menu.
 - Search by using keywords—Click the search icon (magnifying glass), enter the search term in the text box, and press Enter. The search results are displayed on the same page.
 - Filter the data displayed in the table—Click the filter icon (funnel) and select whether you want to show or hide advanced filters. You can then add or remove filter criteria, save criteria as a filter, apply or clear filters, and so on. The filtered results are displayed on the same page.

Sections on the Page

The following topics describe the sections on the Interfaces Health page.

Gauge Chart

The gauge chart displays the overall health of your network devices, as a percentage. Hover over the chart to view the percentage of healthy devices in based on routing health. The percentage of healthy devices based on routing health is also displayed below the chart. The percentage is 0 even if one routing KPI in the network is unhealthy or has an alert associated to it.

List of KPIs

KPIs related to routing components such as BGP, IS-IS, OSPF, and RIB that affect routing health are listed next to the gauge chart. The overall health of a KPI is displayed next to the KPI itself. Click any KPI to view more factors that affect the overall health.

For example, click **LSP** to view other KPIs (such as LSP State and LSF Flap) that contribute to the overall percentage that is displayed. In some instances, you can also view a second level of KPIs that affect the overall percentage. For example, click **LSP** to view other related KPIs that affect the overall health.

A total of four KPIs are listed by default next to the gauge chart. At the end of the KPI list, you can select more KPIs from the **Select Others** drop-down list to view additional KPIs such as MPLS and FIB. You can click any KPI to view more factors that affect the overall health. To clear all selected KPIs, click **Clear Selection**.

Overview Tab

The Overview tab consists of the Devices, OS Versions, and Models widget.

The Devices widget displays the overall percentage of healthy and unhealthy devices (based on routing health) in a donut chart. However, when you select a KPI from the list next to the gauge chart, the Devices widget displays the overall health of the device in relation to that KPI. The total number of devices in the network is displayed in the middle of the chart. Healthy devices are highlighted in green in the donut chart, and unhealthy devices are highlighted in red. Hover over the green portion of the donut chart to view the percentage and the total number of healthy devices. Hover over the red portion of the donut chart to view the percentage and the total number of unhealthy devices.

The OS Versions widget displays the percentage of healthy and unhealthy devices (based on routing health) grouped by OS versions in a stacked bar chart. Hover over the bar chart to view this percentage. The portion of the bar chart highlighted in green represents healthy devices. The portion of the bar chart highlighted in red represents unhealthy devices. A maximum of ten OS versions are displayed in the widget in the descending order of the unhealthy device percentage.

The Models widget displays the percentage of healthy and unhealthy devices (based on routing health) grouped by device models in a stacked bar chart. Hover over the bar chart to view this percentage. The portion of the bar chart highlighted in green represents healthy devices. The portion of the bar chart highlighted in red represents unhealthy devices. A maximum of ten device models are displayed in the widget in the descending order of the unhealthy device percentage.

Timeline Tab

The Timeline tab displays the percentage of healthy devices for a predefined (a week, a day, 3 hours, 1 hour, 30 minutes) or a custom period on a graph. The percentage of healthy devices is calculated based on routing components such as BGP, IS-IS, OSPF, RSVP, LDP neighbors; LSP state and flaps; and the state of RIB and FIB. By default, performance for the past 30 minutes is displayed. To change this period, click the **Week**, **Day**, **3 hrs**, **1 hr**, **30 mins**, or **Custom** buttons provided above the graph.

You can also use the range picker provided above the graph to view information for a predefined or custom period. You can click the arrows provided on either side of the range picker to move between ranges. For example, click the **Week** button to view routing health information for a week. Click the arrow on the left of the range picker to view information for the previous week. Click the arrow on the right of range picker to view information for the following week.

Data lines for the overall routing health and routing health in relation to individual KPIs (BGP, IS-IS, OSPF, RIB, and so on) are displayed on the graph by default. Hover over any point on the line graph to view the percentage health for that particular point. When you select a KPI from the list next to the gauge chart, the graph is updated to only show the overall routing health in relation to that KPI. The KPIs are also provided below the graph with colors automatically assigned to each of the KPIs. Click any KPI to view a data line in the assigned color on the graph. Click **Clear Selection** to clear all selected KPIs.

Affected Items Tab

The Affected Items tab consists of the Devices tab.

The Devices tab displays details of unhealthy devices (based on routing health) in the order of severity of the alerts generated on those devices. The various alerts are:

- **Action Needed**—The routing neighbors and entries in the routing and forwarding tables have issues that you must address.
- **Urgent Action Needed**—The routing neighbors and entries in the routing and forwarding tables have issues that you must address immediately.

Click any device hostname to view alerts for that device on the Alerts for *Device-Hostname* page.

Distribution Tab

The Distribution tab consists of the OS, Device Model, and Sites tabs.

The OS tab displays the details of unhealthy devices (based on routing health) grouped by the device OS version. The details on this tab are updated based on the KPIs that you select. You can view the total number of affected devices, total number devices that use the OS version, and the percentage impact for each OS version on the tab. Click the link in the **Devices Affected** column to view details of all affected devices that use that OS version, on the Devices Affected for *OS-Name* page. Click the link in the **Total Network** column to view details of all devices that use that OS version, on the Total Devices for *OS-Name* page.

The Device Model tab displays the details of unhealthy devices (based on routing health) grouped by device (hardware) model. The details on this tab are updated based on the KPIs that you select. You can view the total number of affected devices, total number of devices that belong to the hardware model category, and the percentage impact for each hardware model category on the tab. Click the link in the **Devices Affected** column to view details of all affected devices that belong to that hardware model category, on the Devices Affected for *Model-Name* page. Click the link in the **Total Network** column to view all devices that belong to that hardware model category, on Total Devices for *Model-Name* page.

The Sites tab displays the details of unhealthy devices (based on routing health) grouped by sites. A site is the location of a device in an organization. The details on this tab are updated based on the KPIs that you select. You can view the total number of affected devices, total number devices in the site, and the

percentage impact for each site on the tab. Click the link in the **Devices Affected** column to view details of all affected devices that belong to that site, on the Devices Affected for *Site-Name* page. Click the link in the **Total Network** column to view details of all devices that belong to that site, on the Total Devices for *Site-Name* page.

RELATED DOCUMENTATION

[About the WAN Health Tab | 399](#)

About the Trust Tab

IN THIS SECTION

- [Tasks You Can Perform | 417](#)

To access this tab, click **Observability > Health > Health Dashboard > Trust**.

Use this tab to view average trust score of the targets in your network and real-time information about compliance, integrity, and vulnerabilities that affect the overall health of devices.

The Trust tab consists of the following:

- Trust Pane
- Vulnerabilities Accordion
- Compliance Accordion
- Integrity Accordion

The Trust pane displays the trust score, trust plan applied, overall trust trend, and a graph that displays the average trust score. Each accordion on this page has three sections. The first section provides a description of the accordion. The second section displays the latest aggregate health percentage and total number of devices. The last section displays a health trend graph representing the average health of devices for the past 30 days.

Tasks You Can Perform

You can perform the following tasks from this page:

- On the Trust pane, you can:
 - View the average trust score (in percentage) of devices within the network, with 100 percent representing complete trust. A trust score is calculated based on the values of a list of factors such as prerequisites, variable, and reputational factors. See ["Trust Score Overview" on page 621](#).

The icon before the trust score indicates how trustworthy a device is.

- A green icon indicates that the average trust score is 90 percent or above.
- An orange icon indicates that the average trust score is between 40 and 90 percent.
- A red icon indicates that the average trust score is 40 percent or below.
- View the predefined trust plan applied to the network.

You can change the current predefined plan from the **Plan** drop-down list on the Trust Settings page.

- View an overall trend (in percentage) of the improvement or decline of the average trust score for the past 30 days.

A green arrow pointing up before the overall trend percentage indicates that the average score has improved. A red arrow pointing down indicates that the average trust score has decreased since the details on the Trust tab was last updated.

- View a graph of the average trust score for the past 30 days. You can also view two lines showing the high (when average trust score is 90 percent and above) and low (when average trust score is below 40 percent) threshold levels of the average trust score.
- View overall device health based on the device vulnerabilities on the Vulnerabilities accordion. See ["About the Vulnerability Health Page" on page 419](#).

On the Vulnerabilities accordion, you can:

- View the average health (in percentage) of devices based on the vulnerabilities on the devices, and the total number of unhealthy devices.
- View the percentage and health of the KPIs, and the total number of unhealthy devices in relation to that KPI.
- View a graph of the average health of devices based on device vulnerabilities for the past 30 days.
- View overall device health based on compliance of devices on the Compliance accordion. See ["About the Compliance Health Page" on page 423](#).

On the Compliance accordion, you can:

- View the average health (in percentage) of devices based on compliance of devices, and the total number of unhealthy devices.
- View the percentage and health of the KPIs, and the total number of unhealthy devices in relation to that KPI.
- View a graph of the average health of devices based on compliance of devices for the past 30 days.
- View overall device health based on integrity of devices on the Integrity accordion. See ["About the Integrity Health Page" on page 427](#).

On the Integrity accordion, you can:

- View the average health (in percentage) of devices based on integrity of devices, and the total number of unhealthy devices.
- View the percentage and health of the KPIs, and the total number of unhealthy devices in relation to that KPI.
- View a graph of the average health of devices based on integrity of devices for the past 30 days.
- Expand or collapse accordions.
 - Click **Expand All** to expand the accordion on the tab. You can view detailed information that is related to KPIs that affect overall health based on alerts, and a performance graph.
 - Click **Collapse All** to collapse all accordions on the tab. You can view Trust score pane and overall health of the accordion that is based on alerts, and a condensed view of the performance graph.
- View the total number of devices in the network in the top-right corner of the tab.
- View the timestamp of the last update or refresh (in *MMM DD, YYYY hh:mm:ss AM/PM* format) in the top-right corner of the tab. For example: May 12, 2024 12:42:43 PM.
- Click the **Settings** (gear) icon to configure these parameters to run a compliance scan on devices in the network:
 - **Plan**—Select a predefined trust plan to calculate the trust score for each device. For more information, see ["Trust Plans Overview" on page 617](#).
 - **Benchmark**—Select a benchmarks document that lists compliance policies and rules defined by the Center for Internet Security (CIS). For more information, see ["About the Compliance Benchmarks Page" on page 608](#).

- **Profile**—Select a benchmark profile depending on the level of security of the scan. Profiles in the benchmarks document defines the rules and policies enforced on devices to obtain a desired level of compliance. For more information, see ["About the Compliance Benchmarks Page" on page 608](#).
- **Tailoring**—Select a tailoring document to assess the compliance of devices. You can create customized tailoring documents for one or more rules defined in the benchmarks document. For more information, see ["About the Compliance Tailorings Page" on page 610](#).

Click **Apply** to save.

You are returned to the Trust tab where you can view results of the scan on the Compliance accordion based on the configured parameters. The Compliance accordion is displayed only when the scan matches the Trust settings that you configured.

- Update the details on the tab by clicking the refresh icon.

RELATED DOCUMENTATION

[About the Vulnerability Health Page | 419](#)

[About the Compliance Health Page | 423](#)

[About the Integrity Health Page | 427](#)

About the Vulnerability Health Page

IN THIS SECTION

● [Tasks You Can Perform | 420](#)

● [Sections on the Page | 421](#)

You can access this page in the following ways:

Click **Observability > Health > Health Dashboard > Trust (Tab) > Vulnerabilities (Accordion) > View Details**.

You can view detailed information about the health of devices based on the vulnerabilities on the devices.

Click **Observability > Health > Health Dashboard > Trust (Tab) > Vulnerabilities (Accordion) > Feature (click any KPI listed)**.

You can view information about the overall health of devices based on the KPI that you selected.

The Vulnerability Health page consists of a gauge chart, a list of KPIs, Overview tab, Timeline tab, Affected Items tab, and Distribution tab.

Tasks You Can Perform

You can perform the following tasks from this page:

- View a gauge chart of the overall health of devices based on device vulnerabilities in your network.
- View a list of KPIs that affect the health of devices. The KPIs listed are Advisories and PBNs.
- On the Overview tab:
 - View the percentage of healthy and unhealthy devices in the donut chart.
 - View the percentage of healthy and unhealthy devices (based on device vulnerabilities) grouped by OS versions in a stacked bar chart.
 - View the percentage of healthy and unhealthy devices (based on device vulnerabilities) grouped by device models in a stacked bar chart.
- On the Timeline tab:
 - View the percentage of healthy devices based on device vulnerabilities for the past 30 days on a graph. The graph is updated based on the KPIs that you select.
- On the Affected Items tab:
 - View the details of unhealthy devices (based on device vulnerabilities) in the order of severity of the alerts generated on those devices.
 - Click any device hostname to view alerts for that device on the Vulnerabilities for *Device-Hostname* page.
- The Distribution tab consists of the OS tab and Device Model tab.
 - On the OS tab, view the details of unhealthy devices (based on device vulnerabilities) grouped by device OS version.
 - On the Device Model tab, view the details of unhealthy devices grouped by device model.
- Click **Clear Selection** to clear all selected KPIs.
- View the total number of devices in the network on the top-right corner of the page.

- View the timestamp of the last update or refresh (in *MMM DD, YYYY hh:mm:ss AM/PM* format) in the top-right corner of the tab. For example: May 24, 2024 12:42:43 PM
- Update the page by clicking the refresh icon.
- On the Affected Items and Distribution tab, you can also:
 - Sort, resize, or re-arrange columns in a table (grid).
 - Show or hide columns in the table or reset page preferences, using the vertical ellipsis menu.
 - Search by using keywords—Click the search icon (magnifying glass), enter the search term in the text box, and press Enter. The search results are displayed on the same page.
 - Filter the data displayed in the table—Click the filter icon (funnel) and select whether you want to show or hide advanced filters. You can then add or remove filter criteria, save criteria as a filter, apply or clear filters, and so on. The filtered results are displayed on the same page.
 - Click the Display drop-down at the bottom of the tab to select the maximum number of entries from the available options to be displayed in the table. The default number of entries displayed is **15**.

Sections on the Page

The following topics describe the sections on the Vulnerability Health page.

Gauge Chart

The gauge chart displays the overall health of network devices as a percentage. Hover over the chart to view the percentage of healthy devices based on device vulnerabilities. The percentage of healthy devices is also displayed below the chart. The percentage is 0 even if one device in the network is unhealthy or has an alert associated with it.

List of KPIs

KPIs such as PBNs and Advisories are listed next to the gauge chart. The overall health of a KPI is displayed next to the KPI itself. A total of two KPIs are listed by default next to the gauge chart. To clear all selected KPIs, click **Clear Selection**.

Overview Tab

The Overview tab consists of the Devices, OS Versions, and Models widgets.

The Devices widget displays the overall percentage of healthy and unhealthy devices (based on device vulnerabilities) in a donut chart. However, when you select a KPI from the list next to the gauge chart,

the Devices widget displays the overall health of the device based on that KPI. Healthy devices are highlighted in green in the donut chart, and unhealthy devices are highlighted in red. Hover over the green portion of the donut chart to view the percentage and total number of healthy devices. Hover over the red portion of the donut chart to view the percentage and total number of unhealthy devices.

The OS Versions widget displays the percentage of healthy and unhealthy devices (based on device vulnerabilities) grouped by OS versions in a stacked bar chart. Hover over the bar chart to view this percentage. The portion of the bar chart highlighted in green represents healthy devices. The portion of the bar chart highlighted in red represents unhealthy devices.

The Models widget displays the percentage of healthy and unhealthy devices (based on device vulnerabilities) that are grouped by device models in a stacked bar chart. Hover over the bar chart to view this percentage. The portion of the bar chart highlighted in green represents healthy devices. The portion of the bar chart highlighted in red represents unhealthy devices.

Timeline Tab

The Timeline tab displays the percentage of healthy devices based on device vulnerabilities for the past 30 days on a graph.

Data lines for the overall device health and device health in relation to individual KPIs (PBNs and Advisories) are displayed on the graph by default. Hover over any point on the line graph to view the percentage health for that particular point. When you select a KPI from the list next to the gauge chart, the graph is updated to only show the overall health of the device in relation to that KPI. The KPIs are also provided below the graph with colors automatically assigned to each of the KPIs. Click any KPI to view a data line in the assigned color on the graph. Click **Clear Selection** to clear all selected KPIs.

Affected Items Tab

The Affected Items tab consists of the Devices tab.

The Devices tab displays details of unhealthy devices in the order of severity of the alerts generated on those devices. Click any device hostname to view alerts for that device on the Vulnerabilities for *Device-Hostname* page.

Distribution Tab

The Distribution tab consists of the OS tab and Device Model tab.

The OS tab displays the details of unhealthy devices (based on device vulnerabilities) grouped by device OS version. The details on this tab are updated based on the KPIs that you select. Click the link in the **Devices Affected** column provided next to the name of the OS version to view details of all affected devices that use that OS version. This information is displayed on the Vulnerabilities for *OS-Name* page. The percentage impact for each OS version is also displayed on the OS tab.

The Device Model tab displays the details of unhealthy devices grouped by device model. The details on this tab are updated based on the KPIs that you select. Click the link in the **Devices Affected** column provided next to the name of the hardware model to view details of all affected devices that belong to that hardware model category. This information is displayed on the Vulnerabilities for *Model-Name* page. The percentage impact for each device model is also displayed on the Device Model tab.

RELATED DOCUMENTATION

[About the Trust Tab | 416](#)

About the Compliance Health Page

IN THIS SECTION

- [Tasks You Can Perform | 423](#)
- [Sections on the Page | 425](#)

You can access this page in the following ways:

- Click **Observability > Health > Health Dashboard > Trust (Tab) > Compliance (Accordion) > View Details**.

You can view detailed information about the health of devices based on device compliance.

- Click **Observability > Health > Health Dashboard > Trust (Tab) > Compliance (Accordion) > Feature (click any KPI listed)**.

You can view information about the overall health of devices based on the KPI that you selected.

The Compliance Health page consists of a gauge chart, a list of KPIs, Overview tab, Timeline tab, Affected Items tab, and Distribution tab.

Tasks You Can Perform

You can perform the following tasks from this page:

- View a gauge chart of the overall health of devices that are based on the compliance of devices in your network.

- View a list of KPIs that affect the health of devices. The KPIs displayed on the page are provided in the benchmark document that is defined by the Center for Internet Security (CIS). You can select more KPIs from the **Select Others** drop down list available at the end of KPIs list.
- On the Overview tab:
 - View the percentage of healthy and unhealthy devices in the donut chart.
 - View the percentage of healthy and unhealthy devices (based on device compliance) grouped by OS versions in a stacked bar chart.
 - View the percentage of healthy and unhealthy devices (based on device compliance) grouped by device models in a stacked bar chart.
- On the Timeline tab:
 - View the percentage of healthy devices based on device compliance for the past 30 days on a graph. The graph is updated based on the KPIs that you select.
- On the Affected Items tab:
 - View the details of unhealthy devices (based on device compliance) in the order of severity of the alerts generated on those devices.
 - Click any device hostname to view alerts for that device on the Rule Results for *Device-Hostname* page.
- The Distribution tab consists of the OS tab and Device Model tab.
 - On the OS tab, view the details of unhealthy devices (based on device compliance) grouped by device OS version.
 - On the Device Model tab, view the details of unhealthy devices grouped by device model.
- Click **Clear Selection** to clear all selected KPIs.
- View the total number of devices in the network on the top-right corner of the page.
- View the timestamp of the last update or refresh (in *MMM DD, YYYY hh:mm:ss AM/PM* format) in the top-right corner of the tab. For example: May 24, 2024 12:42:43 PM
- Update the page by clicking the refresh icon.
- On the Affected Items and Distribution tab, you can also:
 - Sort, resize, or re-arrange columns in a table (grid).
 - Show or hide columns in the table or reset page preferences, using the vertical ellipsis menu.

- Search by using keywords—Click the search icon (magnifying glass), enter the search term in the text box, and press Enter. The search results are displayed on the same page.
- Filter the data displayed in the table—Click the filter icon (funnel) and select whether you want to show or hide advanced filters. You can then add or remove filter criteria, save criteria as a filter, apply or clear filters, and so on. The filtered results are displayed on the same page.
- Click the Display drop-down at the bottom of the tab to select the maximum number of entries from the available options to be displayed in the table. The default number of entries displayed is **15**.

Sections on the Page

The following topics describe the sections on the Compliance Health page.

Gauge Chart

The gauge chart displays the overall health of network devices as a percentage. Hover over the chart to view the percentage of healthy devices based on device compliance. The percentage of healthy devices is also displayed below the chart. The percentage is 0 even if one device in the network is unhealthy or has an alert associated with it.

List of KPIs

KPIs are listed next to the gauge chart. The overall health of a KPI is displayed next to the KPI. The KPIs displayed on the page are provided in the benchmark document that is defined by the Center for Internet Security (CIS). At the end of the KPIs list, you can select more KPIs from the **Select Others** drop-down list to view additional KPIs. To clear all selected KPIs, click **Clear Selection**.

Overview Tab

The Overview tab consists of the Devices, OS Versions, and Models widgets.

The Devices widget displays the overall percentage of healthy and unhealthy devices (based on device compliance) in a donut chart. However, when you select a KPI from the list next to the gauge chart, the Devices widget displays the overall health of the device based on that KPI. Healthy devices are highlighted in green in the donut chart, and unhealthy devices are highlighted in red. Hover over the green portion of the donut chart to view the percentage and total number of healthy devices. Hover over the red portion of the donut chart to view the percentage and total number of unhealthy devices.

The OS Versions widget displays the percentage of healthy and unhealthy devices (based on device compliance) grouped by OS versions in a stacked bar chart. Hover over the bar chart to view this

percentage. The portion of the bar chart highlighted in green represents healthy devices. The portion of the bar chart highlighted in red represents unhealthy devices.

The Models widget displays the percentage of healthy and unhealthy devices (based on device compliance) that are grouped by device models in a stacked bar chart. Hover over the bar chart to view this percentage. The portion of the bar chart highlighted in green represents healthy devices. The portion of the bar chart highlighted in red represents unhealthy devices.

Timeline Tab

The Timeline tab displays the percentage of healthy devices that are based on device compliance for the past 30 days on a graph.

Data lines for the overall device health and device health in relation to individual KPIs (provided in the benchmark document) are displayed on the graph by default. Hover over any point on the line graph to view the percentage health for that particular point. When you select a KPI from the list next to the gauge chart, the graph is updated to only show the overall health of the device in relation to that KPI. The KPIs are also provided below the graph with colors automatically assigned to each of the KPIs. Click any KPI to view a data line in the assigned color on the graph. Click **Clear Selection** to clear all selected KPIs.

Affected Items Tab

The Affected Items tab consists of the Devices tab.

The Devices tab displays details of unhealthy devices in the order of severity of the alerts that are generated on those devices. Click any device hostname to view alerts for that device on the Rule Results for *Device-Hostname* page.

Distribution Tab

The Distribution tab consists of the OS tab and Device Model tab.

The OS tab displays the details of devices (based on device compliance) grouped by device OS version. The details include the number of devices and the average trust score of those devices. The details on this tab are updated based on the KPIs that you select. The percentage impact for each OS version is also displayed on the OS tab.

The Device Model tab displays the details of devices that are grouped by device model. The details include the number of devices and the average trust score of those devices. The details on this tab are updated based on the KPIs that you select. The percentage impact for each device model is also displayed on the Device Model tab.

About the Integrity Health Page

IN THIS SECTION

- [Tasks You Can Perform | 427](#)
- [Sections on the Page | 428](#)

You can access this page in the following ways:

- Click **Observability > Health > Health Dashboard > Trust (Tab) > Integrity (Accordion) View Details**.

You can view detailed information about the health of devices based on integrity of devices.

- Click **Observability > Health > Health Dashboard > Trust (Tab) > Integrity (Accordion) > Feature (click any KPI listed)**.

You can view information about the overall health of devices based on the KPI that you selected.

The Integrity Health page consists of a gauge chart, a list of KPIs, Overview tab, Timeline tab, Affected Items tab, and Distribution tab.

Tasks You Can Perform

You can perform the following tasks from this page:

- View a gauge chart of the overall health of devices based on integrity of devices in your network.
- View a list of KPIs that affect the health of devices. The KPIs listed are Hardware EOL and Software EOL.
- On the Overview tab:
 - View the percentage of healthy and unhealthy devices in the donut chart.
 - View the percentage of healthy and unhealthy devices (based on integrity of devices) grouped by OS versions in a stacked bar chart.
 - View the percentage of healthy and unhealthy devices (based on integrity of devices) grouped by device models in a stacked bar chart.
- On the Timeline tab:

- View the percentage of healthy devices based on integrity of devices for the past 30 days on a graph. The graph is updated based on the KPIs that you select.
- On the Affected Items tab:
 - View the details of unhealthy devices (based on integrity of devices) in the order of severity of the alerts generated on those devices.
 - Click any device hostname to view alerts for that device on the Integrity page.
- The Distribution tab consists of the OS tab and Device Model tab.
 - On the OS tab, view the details of unhealthy devices (based on integrity of devices) grouped by device OS version.
 - On the Device Model tab, view the details of unhealthy devices (based on integrity of devices) grouped by device model.
- Click **Clear Selection** to clear all selected KPIs.
- View the total number of devices in the network on the top-right corner of the page.
- View the timestamp of the last update or refresh (in *MMM DD, YYYY hh:mm:ss AM/PM* format) in the top-right corner of the tab. For example: May 24, 2024 12:42:43 PM
- Update the page by clicking the refresh icon.
- On the Affected Items and Distribution tab, you can also:
 - Sort, resize, or re-arrange columns in a table (grid).
 - Show or hide columns in the table or reset page preferences, using the vertical ellipsis menu.
 - Search by using keywords—Click the search icon (magnifying glass), enter the search term in the text box, and press Enter. The search results are displayed on the same page.
 - Filter the data displayed in the table—Click the filter icon (funnel) and select whether you want to show or hide advanced filters. You can then add or remove filter criteria, save criteria as a filter, apply or clear filters, and so on. The filtered results are displayed on the same page.
 - Click the Display drop-down at the bottom of the tab to select the maximum number of entries from the available options to be displayed in the table. The default number of entries displayed is **15**.

Sections on the Page

The following topics describe the sections on the Integrity Health page.

Gauge Chart

The gauge chart displays the overall health of network devices as a percentage. Hover over the chart to view the percentage of healthy devices based on integrity of devices. The percentage of healthy devices is also displayed below the chart. The percentage is 0 even if one device in the network is unhealthy or has an alert associated with it.

List of KPIs

KPIs such as Hardware EOL and Software EOL are listed next to the gauge chart. The overall health of a KPI is displayed next to the KPI itself. To clear all selected KPIs, click **Clear Selection**.

Overview Tab

The Overview tab consists of the Devices, OS Versions, and Models widgets.

The Devices widget displays the overall percentage of healthy and unhealthy devices (based on integrity of devices) in a donut chart. However, when you select a KPI from the list next to the gauge chart, the Devices widget displays the overall health of the device based on that KPI. Healthy devices are highlighted in green in the donut chart, and unhealthy devices are highlighted in red. Hover over the green portion of the donut chart to view the percentage and total number of healthy devices. Hover over the red portion of the donut chart to view the percentage and total number of unhealthy devices.

The OS Versions widget displays the percentage of healthy and unhealthy devices (based on integrity of devices) grouped by OS versions in a stacked bar chart. Hover over the bar chart to view this percentage. The portion of the bar chart highlighted in green represents healthy devices. The portion of the bar chart highlighted in red represents unhealthy devices.

The Models widget displays the percentage of healthy and unhealthy devices (based on integrity of devices) that are grouped by device models in a stacked bar chart. Hover over the bar chart to view this percentage. The portion of the bar chart highlighted in green represents healthy devices. The portion of the bar chart highlighted in red represents unhealthy devices.

Timeline Tab

The Timeline tab displays the percentage of healthy devices based on integrity of devices for the past 30 days on a graph.

Data lines for the overall device health and device health in relation to individual KPIs (Hardware EOL and Software EOL) are displayed on the graph by default. Hover over any point on the line graph to view the percentage health for that particular point. When you select a KPI from the list next to the gauge chart, the graph is updated to only show the overall health of the device in relation to that KPI. The KPIs are also provided below the graph with colors automatically assigned to each of the KPIs. Click any KPI to view a data line in the assigned color on the graph. Click **Clear Selection** to clear all selected KPIs.

Affected Items Tab

The Affected Items tab consists of the Devices tab.

The Devices tab displays details of unhealthy devices in the order of severity of the alerts generated on those devices. Click any device hostname to view alerts for that device on the Integrity page.

Distribution Tab

The Distribution tab consists of the OS tab and Device Model tab.

The OS tab displays the details of unhealthy devices (based on integrity of devices) grouped by device OS version. The details on this tab are updated based on the KPIs that you select. Click the link in the **Devices Affected** column provided next to the name of the OS version to view details of all affected devices that use that OS version. This information is displayed on the Integrity - Software EOL page. The end of life (EOL) and end of support (EOS) date for each OS version is also displayed on the OS tab.

The Device Model tab displays the details of unhealthy devices grouped by device model. The details on this tab are updated based on the KPIs that you select. Click the link in the **Devices Affected** column provided next to the name of the device model to view details of all affected devices that belong to that hardware model category. This information is displayed on the Integrity - Hardware EOL page. The EOL and EOS date for each device model is also displayed on the Device Model tab.

Update Trust KPIs in an Air-Gapped Installation

IN THIS SECTION

- [Update Trust KPIs using Paragon Shell | 431](#)

Overview

Routing Director enables you to update Trust-related KPIs such as PBNs, Advisories, and so on to the latest version when your on-premise installation does not have access to the Internet. You can update Trust-related KPIs using the request `paragon trust data update hostname hostname` Paragon Shell CLI command.

After you update the Trust-related KPIs, you can view the updated information on the Trust tab of the Health Dashboard (**Observability > Health > Health Dashboard > Trust**). The Trust tab displays the

percentage and health of Trust-related KPIs on different accordions. Use this information to detect any changes in the Trust score and perform corrective actions, when necessary.

You must run this command only in a maintenance window to ensure there are no disruptions to network operation.

Update Trust KPIs using Paragon Shell

Perform the following steps to update the trust-related KPIs when your on-premise installation does not have access to the Internet:

1. Log in to a controller cluster node as the root user.

You are logged in to Paragon Shell.

2. Type `exit` to exit to the Linux root shell.
3. Create a `/data/trust` folder on the node at the root level.
4. Securely copy the `sirts.json`, `ascSWEol.json`, `ascHWEol.json`, and `pbn.json.gz` files to the `/data/trust` folder on the node.
5. Identify the hostname of the node on which you copied the files, using the `kubectl get nodes -o wide --show-labels | awk -v OFS=' \t\t' '{print $1, $6, $13}'` command.
Map the IP address of the node on which the files are located to the hostname.

```
root@hostname1:~# kubectl get nodes -o wide --show-labels | awk -v OFS=' \t\t' '{print $1, $6, $13}'
```

NAME	INTERNAL-IP	
hostname1	10.1.2.3	beta.kubernetes.io/arch=amd64,beta.kubernetes.io/instance-type=rke2,beta.kubernetes.io/os=linux,kubernetes.io/arch=amd64,kubernetes.io/hostname=hostname1,kubernetes.io/os=linux,node-role.kubernetes.io/control-plane=true,node-role.kubernetes.io/etcd=true,node-role.kubernetes.io/master=true,node.kubernetes.io/instance-type=rke2,role=master
hostname2	10.1.3.4	beta.kubernetes.io/arch=amd64,beta.kubernetes.io/instance-type=rke2,beta.kubernetes.io/os=linux,kubernetes.io/arch=amd64,kubernetes.io/hostname=hostname2,kubernetes.io/os=linux,node-role.kubernetes.io/control-plane=true,node-role.kubernetes.io/etcd=true,node-role.kubernetes.io/master=true,node.kubernetes.io/instance-type=rke2,role=master

<Output Snipped>

The hostname is displayed in the `kubernetes.io/hostname=hostname` format. For example, in the sample output, `hostname1` is the hostname with IP address 10.1.2.3.

6. Type `cli` to re-enter Paragon Shell.
7. Update Trust-related KPIs using the request `paragon trust data update hostname hostname` command.

Where, *hostname* is the hostname of the node where the copied files are located.

This command shuts down all the Trust pods and then starts a new Trust pod to update the data. Once the update process is completed, the command restarts Trust pods in the run mode and automatically resets their default replica values. The default replica values are `minReplicas:2` and `maxReplicas:5`.

For more information, see ["request paragon trust data update hostname" on page 1438](#).

8. (Optional) Monitor the status of the update job using the request `paragon debug logs namespace trust service network` command.

The command generates a `logs-network-date_time.tar.gz` log file. This log file is saved in the `/root/troubleshooting/log` directory.



NOTE:

The logs are available only for a short period after the Trust pods have been restarted.

The following log indicates a successful update:

```
{ "caller": "cmd:serve.go:1144", "level": "info", "message": "Update file for sirts.json found.", "timestamp": "1737990749063", "ts": "2025-01-27T15:12:29.063465117Z" }

{ "caller": "cmd:serve.go:1160", "level": "info", "message": "Update file for pbn.json.gz found.", "timestamp": "1737990749076", "ts": "2025-01-27T15:12:29.076234873Z" }

{ "caller": "cmd:serve.go:1029", "level": "info", "message": "Update file for ascSwEol.json found.", "timestamp": "1737990752433", "ts": "2025-01-27T15:12:32.433507543Z" }

{ "caller": "cmd:serve.go:1049", "level": "info", "message": "Update file for ascHwEol.json found.", "timestamp": "1737990752458", "ts": "2025-01-27T15:12:32.458457932Z" }
```

The following log indicates a partial update:

```
{ "caller": "cmd:serve.go:1740", "level": "info", "message": "Preload Checksum for sirts has not
changed. No Update.", "org_id": "76c5b9b7-0ad9-4614-ad8d-
baddfb46108a", "timestamp": "1737992169212", "ts": "2025-01-27T15:36:09.212820833Z" }

{ "caller": "cmd:serve.go:1740", "level": "info", "message": "Preload Checksum for pbns has not
changed. No Update.", "org_id": "76c5b9b7-0ad9-4614-ad8d-
baddfb46108a", "timestamp": "1737992169215", "ts": "2025-01-27T15:36:09.215643375Z" }

{ "caller": "cmd:serve.go:1744", "id": "sweol", "level": "info", "message": "Preloading
metadata", "org_id": "76c5b9b7-0ad9-4614-ad8d-
baddfb46108a", "timestamp": "1737992169216", "ts": "2025-01-27T15:36:09.216304717Z" }

{ "caller": "cmd:serve.go:1744", "id": "hweol", "level": "info", "message": "Preloading
metadata", "org_id": "76c5b9b7-0ad9-4614-ad8d-
baddfb46108a", "timestamp": "1737992169394", "ts": "2025-01-27T15:36:09.394893267Z" }
```

The Trust database is automatically updated when Routing Director detects new or changed information.

9. Check the Trust tab on the Health Dashboard (**Observability > Health > Health Dashboard > Trust**) of the Routing Director GUI to verify the updates.
10. (Optional) To troubleshoot any errors that may occur when you update Trust-related KPIs, contact the [Juniper Networks Technical Assistance Center \(JTAC\)](#) team.

SEE ALSO

[request paragon trust data update hostname](#) | 1438

About the Active Assurance Tab

IN THIS SECTION

- [Task You Can Perform](#) | 434

To access this tab, click **Observability > Health > Health Dashboard > Active Assurance**.

This tab provides a unified view of the overall health of various measurements in your network in real-time. You can monitor the health of various measurements, and use this data to perform corrective actions when necessary.

The Active Assurance tab consists of the following accordions:

- **Applications**—Use this accordion to view the overall health of measurements related to HTTP, DNS, and RPM HTTP plug-ins in your network. For more information, see ["About the Applications Page" on page 435](#).
- **L3 Network**—Use this accordion to view the overall health of measurements related to TCP, UDP, TWAMP, TWAMP Reflector, Ping, RPM TCP, RPM UDP, RPM TWAMP, and RPM Ping plug-ins in your network. For more information, see ["About the L3 Network Page" on page 439](#).
- **L2 Network**—Use this accordion to view the overall health of measurements related to Y.1731 SLM and Y.1731 DM plug-ins in your network. For more information, see ["About the L2 Network Page" on page 442](#).
- **TV & Multicast**—Use this accordion to view the overall health of measurements related to IPTV, OTT, and Netflix Speedtest plug-ins in your network. For more information, see ["About the TV & Multicast Page" on page 446](#).

Each accordion has three sections. The first section provides a description of the accordion. The second section displays the latest aggregate health snapshot. The last section displays a health trend graph representing the average health for the past 30 minutes.

Task You Can Perform

You can perform the following tasks from this page:

- View information related to the measurement health of the Applications, L3 Network, L2 Network, and TV & Multicast accordions.

On the accordions, you can:

- View a health snapshot of the average health (in percentage) of measurements, and the total number of unhealthy measurements. The data displayed for healthy measurements is calculated based on the number of alerts that the measurement has generated.
- View a list of plug-ins and their average health (in percentage) of measurements, and the total number of unhealthy measurements. You can also click any plug-in (listed under Feature) to view more information about the performance of that plug-in.
- View a graph of the average health trend of the measurements for the past 30 minutes. Hover over any point in the graph to view the average health score for that category at that time.
- Expand or collapse accordions.
 - Click **Expand All** to expand the accordions on the tab. You can view detailed information related to KPIs that affect overall health based on alerts, and a performance graph.
 - Click **Collapse All** to collapse all accordions on the tab. You can only view overall health that is based on alerts, and a performance graph.
- View the total number of devices in the network in the top-right corner of the tab.
- View the timestamp of the last update or refresh (in *DD/MM/YYYY hh:mm:ss AM/PM* format) in the top-right corner of the tab. For Example: 27/11/2024 10:32:23 PM.
- Update the details on the tab by clicking the refresh icon.

About the Applications Page

IN THIS SECTION

- [Tasks You Can Perform | 436](#)
- [Sections on the Page | 437](#)

You can access the Applications page in the following ways:

- Click **Observability > Health > Health Dashboard > Active Assurance (Tab) > Applications (Accordion) > View Details**.
- Click **Observability > Health > Health Dashboard > Active Assurance (Tab) > Applications (Accordion) > Plug-ins (click any KPI listed)**.

You can view detailed information on the overall health of measurements related to HTTP, DNS, and RPM HTTP plug-ins in your network.

The Applications page consists of a gauge chart, list of plug-ins and their KPIs, Overview tab, Timeline tab, Affected Items tab, and Distribution tab.

Tasks You Can Perform

You can perform the following tasks from this page:

- View a gauge chart that displays the overall health of various measurements.
- View a list of plug-ins (HTTP, DNS, and RPM HTTP) and their KPIs that affect measurement health.
- On the Overview tab:
 - View the percentage and the total number of healthy and unhealthy measurements in the donut chart. The total number of measurements in the network is displayed in the middle of the donut chart.
- On the Timeline tab:
 - View the percentage of healthy measurements for a predefined or a custom period on a graph. The graph is updated based on the KPIs that you select.
- On the Affected Items tab:
 - View the details of Measurements, Test Agents, and Devices affected by the unhealthy measurements in the order of severity of the alerts generated.
- The Distribution tab consists of the Metrics tab and Sites tab.
 - On the Measurements tab, view the details of unhealthy measurements grouped by plug-ins.
 - On the Sites tab, view the details of unhealthy measurements grouped by sites.
- Click **Clear Selection** to clear all selected KPIs.
- View the total number of measurements in the network on the top-right corner of the page.
- View the time stamp of the last update or refresh (in *MMM DD, YYYY hh:mm:ss AM/PM* format) in the top-right corner of the tab. For Example: May 12, 2024 12:42:43 PM
- Update the page by clicking the refresh icon.
- On the Affected Items and Distribution tab, you can also:
 - Sort, resize, or re-arrange columns in a table (grid).

- Show or hide columns in the table or reset page preferences, using the vertical ellipsis menu.
- Search by using keywords—Click the search icon (magnifying glass), enter the search term in the text box, and press Enter. The search results are displayed on the same page.
- Filter the data displayed in the table—Click the filter icon (funnel) and select whether you want to show or hide advanced filters. You can then add or remove filter criteria, save criteria as a filter, apply or clear filters, and so on. The filtered results are displayed on the same page.

Sections on the Page

The following topics describe the sections on the Applications page.

Gauge Chart

The gauge chart displays the overall health of measurements in terms of percentage. Hover over the chart to view the percentage of healthy measurements in relation to their overall health. The percentage of healthy measurements is also displayed below the chart. The percentage is 0 even if one measurement in the network is unhealthy or has an alert associated to it.

List of Plug-ins and KPIs

The plug-ins, HTTP, DNS, and RPM HTTP, are listed next to the gauge chart. The overall health of a plug-in is displayed next to the plug-in itself. Click any plug-in to view more KPIs that affect the overall health of that plug-in. For more information on plug-ins, see [Table 263 on page 947](#).

For example, click **HTTP** to view other KPIs (such as Events, Metrics, and No data) that contribute to the overall percentage that is displayed. You can also view a second level and third level KPIs that affect the overall measurement health. For example, click the KPI **Metrics** to view metrics like connect time and minimum response time related to the plug-in you selected, which are another related KPIs that affect the overall measurement health.

A total of three plug-ins are listed by default next to the gauge chart. To clear all selected KPIs, click **Clear Selection**.

Overview Tab

The Overview tab has the Measurements widget.

The Measurements widget displays the overall percentage of healthy and unhealthy measurements in a donut chart. However, when you select a plug-in from the list next to the gauge chart, the Measurements widget displays the overall health of the measurements in relation to that plug-in. The total number of measurements in the network is displayed in the middle of the chart. Healthy measurements are highlighted in green in the donut chart, and unhealthy measurements are highlighted

in red. Hover over the green portion of the donut chart to view the percentage and total number of healthy measurements. Hover over the red portion of the donut chart to view the percentage and total number of unhealthy measurements.

Timeline Tab

The Timeline tab displays the percentage of healthy measurements for a predefined (a week, a day, 3 hours, 1 hour, 30 minutes) or a custom period on a graph. By default, performance for the past 30 minutes is displayed. To change this period, click the **Week**, **Day**, **3 hrs**, **1 hr**, **30 mins**, or **Custom** buttons provided above the graph.

You can also use the range picker provided above the graph to view information for a predefined or custom period. You can click the arrows provided on either side of the range picker to move between ranges. For example, click the **Week** button to view device health information for a week. Click the arrow on the left of the range picker to view information for the previous week. Click the arrow on the right of range picker to view information for the following week.

Data lines for the overall measurements and measurements in relation to individual plug-ins are displayed on the graph by default. Hover over any point on the line graph to view the percentage health for that particular point. When you select a plug-in from the list next to the gauge chart, the graph is updated to only show the overall health of the device in relation to that plug-in. Click any KPI to view a data line in the assigned color on the graph. Click **Clear Selection** to clear all selected KPIs.

Affected Items Tab

The Affected Items tab consists of the Measurements, Test Agents, and Devices tabs.

- The Measurements tab displays details of unhealthy measurements in the order of severity of the alerts generated. The various alerts are Action Needed and Urgent Action needed.
- The Test Agents tab displays details of Test Agents associated with the unhealthy measurements in the order of severity of the alerts generated. The various alerts are Action Needed and Urgent Action needed.
- The Devices tab displays details of devices associated with the unhealthy measurement in the order of severity of the alerts generated. The various alerts are Action Needed and Urgent Action needed.

Distribution Tab

The Distribution tab consists of the Metrics and Sites tabs.

The Metrics tab displays the list and details of unhealthy measurements grouped by metrics. The details on this tab are updated based on the KPIs that you select. You can view details such as name, affected measurements, total measurement, and health.

The Sites tab displays the details of unhealthy measurements grouped by sites. A site is the location of a device in an organization. The details on this tab are updated based on the KPIs that you select. You can view the name and the total number of affected measurements.

About the L3 Network Page

IN THIS SECTION

- [Tasks You Can Perform | 439](#)
- [Sections on the Page | 440](#)

You can access the L3 Network page in the following ways:

- Click **Observability > Health > Health Dashboard > Active Assurance (Tab) > L3 Network (Accordion) > View Details**.
- Click **Observability > Health > Health Dashboard > Active Assurance (Tab) > L3 Network (Accordion) > *Plug-ins*** (click any KPI listed).

You can view detailed information on the overall health of measurements related to TCP, UDP, TWAMP, TWAMP Reflector, Ping, RPM TCP, RPM UDP, RPM TWAMP, and RPM Ping plug-ins in your network.

The L3 Network page consists of a gauge chart, list of plug-ins and their KPIs, Overview tab, Timeline tab, Affected Items tab, and Distribution tab.

Tasks You Can Perform

You can perform the following tasks from this page:

- View a gauge chart that displays the overall health of various measurements.
- View a list of plug-ins (TCP, UDP, TWAMP, TWAMP Reflector, Ping, RPM TCP, RPM UDP, RPM TWAMP, and RPM Ping) and their KPIs that affect measurement health.
- On the Overview tab:
 - View the percentage and the total number of healthy and unhealthy measurements in the donut chart. The total number of measurements in the network is displayed in the middle of the donut chart.

- On the Timeline tab:
 - View the percentage of healthy measurements for a predefined or a custom period on a graph. The graph is updated based on the KPIs that you select.
- On the Affected Items tab:
 - View the details of Measurements, Test Agents, and Devices affected by the unhealthy measurements in the order of severity of the alerts generated.
- The Distribution tab consists of the Metrics tab and Sites tab.
 - On the Measurements tab, view the details of unhealthy measurements grouped by plug-ins.
 - On the Sites tab, view the details of unhealthy measurements grouped by sites.
- Click **Clear Selection** to clear all selected KPIs.
- View the total number of measurements in the network on the top-right corner of the page.
- View the time stamp of the last update or refresh (in *MMM DD, YYYY hh:mm:ss AM/PM* format) in the top-right corner of the tab. For Example: May 12, 2024 12:42:43 PM
- Update the page by clicking the refresh icon.
- On the Affected Items and Distribution tab, you can also:
 - Sort, resize, or re-arrange columns in a table (grid).
 - Show or hide columns in the table or reset page preferences, using the vertical ellipsis menu.
 - Search by using keywords—Click the search icon (magnifying glass), enter the search term in the text box, and press Enter. The search results are displayed on the same page.
 - Filter the data displayed in the table—Click the filter icon (funnel) and select whether you want to show or hide advanced filters. You can then add or remove filter criteria, save criteria as a filter, apply or clear filters, and so on. The filtered results are displayed on the same page.

Sections on the Page

The following topics describe the sections on the L3 Network page.

Gauge Chart

The gauge chart displays the overall health of measurements in terms of percentage. Hover over the chart to view the percentage of healthy measurements in relation to their overall health. The percentage of healthy measurements is also displayed below the chart. The percentage is 0 even if one measurement in the network is unhealthy or has an alert associated to it.

List of Plug-ins and KPIs

The plug-ins, TCP, UDP, TWAMP, TWAMP Reflector, Ping, RPM TCP, RPM UDP, RPM TWAMP, and RPM Ping, are listed next to the gauge chart. The overall health of a plug-in is displayed next to the plug-in itself. Click any plug-in to view more KPIs that affect the overall health of that plug-in. For more information on plug-ins, see [Table 263 on page 947](#).

For example, click **TCP** to view other KPIs (such as Events, Metrics, and No data) that contribute to the overall percentage that is displayed. You can also view a second level and third level KPIs that affect the overall measurement health. For example, click the KPI **Metrics** to view metrics such as rate and errored-seconds related to the plug-in you selected, which are another related KPIs that affect the overall measurement health.

A total of four plug-ins are listed by default next to the gauge chart. To clear all selected KPIs, click **Clear Selection**.

Overview Tab

The Overview tab has the Measurements widget.

The Measurements widget displays the overall percentage of healthy and unhealthy measurements in a donut chart. However, when you select a plug-in from the list next to the gauge chart, the Measurements widget displays the overall health of the measurements in relation to that plug-in. The total number of measurements in the network is displayed in the middle of the chart. Healthy measurements are highlighted in green in the donut chart, and unhealthy measurements are highlighted in red. Hover over the green portion of the donut chart to view the percentage and total number of healthy measurements. Hover over the red portion of the donut chart to view the percentage and total number of unhealthy measurements.

Timeline Tab

The Timeline tab displays the percentage of healthy measurements for a predefined (a week, a day, 3 hours, 1 hour, 30 minutes) or a custom period on a graph. By default, performance for the past 30 minutes is displayed. To change this period, click the **Week**, **Day**, **3 hrs**, **1 hr**, **30 mins**, or **Custom** buttons provided above the graph.

You can also use the range picker provided above the graph to view information for a predefined or custom period. You can click the arrows provided on either side of the range picker to move between ranges. For example, click the **Week** button to view device health information for a week. Click the arrow on the left of the range picker to view information for the previous week. Click the arrow on the right of range picker to view information for the following week.

Data lines for the overall measurements and measurements in relation to individual plug-ins are displayed on the graph by default. Hover over any point on the line graph to view the percentage health

for that particular point. When you select a plug-in from the list next to the gauge chart, the graph is updated to only show the overall health of the device in relation to that plug-in. Click any KPI to view a data line in the assigned color on the graph. Click **Clear Selection** to clear all selected KPIs.

Affected Items Tab

The Affected Items tab consists of the Measurements, Test Agents, and Devices tabs.

- The Measurements tab displays details of unhealthy measurements in the order of severity of the alerts generated. The various alerts are Action Needed and Urgent Action needed.
- The Test Agents tab displays details of Test Agents associated with the unhealthy measurements in the order of severity of the alerts generated. The various alerts are Action Needed and Urgent Action needed.
- The Devices tab displays details of devices associated with the unhealthy measurement in the order of severity of the alerts generated. The various alerts are Action Needed and Urgent Action needed.

Distribution Tab

The Distribution tab consists of the Metrics and Sites tabs.

The Metrics tab displays the list and details of unhealthy measurements grouped by metrics. The details on this tab are updated based on the KPIs that you select. You can view details such as name, affected measurements, total measurement, and health.

The Sites tab displays the details of unhealthy measurements grouped by sites. A site is the location of a device in an organization. The details on this tab are updated based on the KPIs that you select. You can view the name and the total number of affected measurements.

About the L2 Network Page

IN THIS SECTION

- [Tasks You Can Perform | 443](#)
- [Sections on the Page | 444](#)

You can access L2 Network page in the following ways:

- Click **Observability > Health > Health Dashboard > Active Assurance (Tab) > L2 Network (Accordion) > View Details**.
- Click **Observability > Health > Health Dashboard > Active Assurance (Tab) > L2 Network (Accordion) > Plug-ins (click any KPI listed)**.

You can view detailed information on the overall health of measurements related to Y.1731 DM and Y.1731 SLM plug-ins in your network.

The L2 Network page consists of a gauge chart, list of plug-ins and its KPIs, Overview tab, Timeline tab, Affected Items tab, and Distribution tab.

Tasks You Can Perform

You can perform the following tasks from this page:

- View a gauge chart that displays the overall health of various measurements.
- View a list of plug-ins (Y.1731 DM and Y.1731 SLM) and their KPIs that affect measurement health.
- On the Overview tab:
 - View the percentage and total number of healthy and unhealthy measurements in the donut chart. The total number of measurements in the network is displayed in the middle of the donut chart.
- On the Timeline tab:
 - View the percentage of healthy measurements for a predefined or a custom period on a graph. The graph is updated based on the KPIs that you select.
- On the Affected Items tab:
 - View the details of Measurements, Test Agents, and Devices affected by the unhealthy measurements in the order of severity of the alerts generated.
- The Distribution tab consists of the Metrics tab and Sites tab.
 - On the Measurements tab, view the details of unhealthy measurements grouped by plug-ins.
 - On the Sites tab, view the details of unhealthy measurements grouped by sites.
- Click **Clear Selection** to clear all selected KPIs.
- View the total number of measurements in the network on the top-right corner of the page.
- View the time stamp of the last update or refresh (in *MMM DD, YYYY hh:mm:ss AM/PM* format) in the top-right corner of the tab. For Example: May 12, 2024 12:42:43 PM

- Update the page by clicking the refresh icon.
- On the Affected Items and Distribution tab, you can also:
 - Sort, resize, or re-arrange columns in a table (grid).
 - Show or hide columns in the table or reset page preferences, using the vertical ellipsis menu.
 - Search by using keywords—Click the search icon (magnifying glass), enter the search term in the text box, and press Enter. The search results are displayed on the same page.
 - Filter the data displayed in the table—Click the filter icon (funnel) and select whether you want to show or hide advanced filters. You can then add or remove filter criteria, save criteria as a filter, apply or clear filters, and so on. The filtered results are displayed on the same page.

Sections on the Page

The following topics describe the sections on the L2 Network page.

Gauge Chart

The gauge chart displays the overall health of measurements in terms of percentage. Hover over the chart to view the percentage of healthy measurements in relation to their overall health. The percentage of healthy measurements is also displayed below the chart. The percentage is 0 even if one measurement in the network is unhealthy or has an alert associated to it.

List of Plug-ins and KPIs

The plug-ins Y.1731 DM and Y.1731 SLM are listed next to the gauge chart. The overall health of a plug-in is displayed next to the plug-in itself. Click any plug-in to view more KPIs that affect the overall health of that plug-in. For more information on plug-ins, see [Table 263 on page 947](#).

For example, click **Y.1731 DM** to view other KPIs (such as Events, Metrics, and No data) that contribute to the overall percentage that is displayed. You can also view a second level and third level KPIs that affect the overall measurement health. For example, click the KPI **Metrics** to view metrics like Two-way delay frames and One-way delay frames that are related to the plug-in you selected and these are another related KPIs that affect the overall measurement health.

A total of two plug-ins are listed by default next to the gauge chart. To clear all selected KPIs, click **Clear Selection**.

Overview Tab

The Overview tab has the Measurements widget.

The Measurements widget displays the overall percentage of healthy and unhealthy measurements in a donut chart. However, when you select a plug-in from the list next to the gauge chart, the Measurements widget displays the overall health of the measurements in relation to that plug-in. The total number of measurements in the network is displayed in the middle of the chart. Healthy measurements are highlighted in green in the donut chart, and unhealthy measurements are highlighted in red. Hover over the green portion of the donut chart to view the percentage and total number of healthy measurements. Hover over the red portion of the donut chart to view the percentage and total number of unhealthy measurements.

Timeline Tab

The Timeline tab displays the percentage of healthy measurements for a predefined (a week, a day, 3 hours, 1 hour, 30 minutes) or a custom period on a graph. By default, performance for the past 30 minutes is displayed. To change this period, click the **Week**, **Day**, **3 hrs**, **1 hr**, **30 mins**, or **Custom** buttons provided above the graph.

You can also use the range picker provided above the graph to view information for a predefined or custom period. You can click the arrows provided on either side of the range picker to move between ranges. For example, click the **Week** button to view device health information for a week. Click the arrow on the left of the range picker to view information for the previous week. Click the arrow on the right of range picker to view information for the following week.

Data lines for the overall measurements and measurements in relation to individual plug-ins are displayed on the graph by default. Hover over any point on the line graph to view the percentage health for that particular point. When you select a plug-in from the list next to the gauge chart, the graph is updated to only show the overall health of the device in relation to that plug-in. Click any KPI to view a data line in the assigned color on the graph. Click **Clear Selection** to clear all selected KPIs.

Affected Items Tab

The Affected Items tab consists of the Measurements, Test Agents, and Devices tabs.

- The Measurements tab displays details of unhealthy measurements in the order of severity of the alerts generated. The various alerts are Action Needed and Urgent Action needed.
- The Test Agents tab displays details of Test Agents associated with the unhealthy measurements in the order of severity of the alerts generated. The various alerts are Action Needed and Urgent Action needed.
- The Devices tab displays details of devices associated with the unhealthy measurement in the order of severity of the alerts generated. The various alerts are Action Needed and Urgent Action needed.

Distribution Tab

The Distribution tab consists of the Metrics and Sites tabs.

The Metrics tab displays the list and details of unhealthy measurements grouped by metrics. The details on this tab are updated based on the KPIs that you select. You can view details such as name, affected measurements, total measurement, and health.

The Sites tab displays the details of unhealthy measurements grouped by sites. A site is the location of a device in an organization. The details on this tab are updated based on the KPIs that you select. You can view the name and the total number of affected measurements.

About the TV & Multicast Page

IN THIS SECTION

- [Tasks You Can Perform | 446](#)
- [Sections on the Page | 447](#)

You can access TV & Multicast page in the following ways:

- Click **Observability > Health > Health Dashboard > Active Assurance (Tab) > TV & Multicast (Accordion) > View Details**.
- Click **Observability > Health > Health Dashboard > Active Assurance (Tab) > TV & Multicast (Accordion) > Plug-ins** (click any KPI listed).

You can view detailed information on the overall health of measurements related to IPTV, OTT, Netflix Speedtest plug-ins in your network.

The TV & Multicast page consists of a gauge chart, list of plug-ins and their KPIs, Overview tab, Timeline tab, Affected Items tab, and Distribution tab.

Tasks You Can Perform

You can perform the following tasks from this page:

- View a gauge chart that displays the overall health of various measurements.
- View a list of plug-ins (IPTV, OTT, Netflix Speedtest) and their KPIs that affect measurement health.

- On the Overview tab:
 - View the percentage and the total number of healthy and unhealthy measurements in the donut chart. The total number of measurements in the network is displayed in the middle of the donut chart.
- On the Timeline tab:
 - View the percentage of healthy measurements for a predefined or a custom period on a graph. The graph is updated based on the KPIs that you select.
- On the Affected Items tab:
 - View the details of Measurements, Test Agents, and Devices affected by the unhealthy measurements in the order of severity of the alerts generated.
- The Distribution tab consists of the Metrics tab and Sites tab.
 - On the Measurements tab, view the details of unhealthy measurements grouped by plug-ins.
 - On the Sites tab, view the details of unhealthy measurements grouped by sites.
- Click **Clear Selection** to clear all selected KPIs.
- View the total number of measurements in the network on the top-right corner of the page.
- View the time stamp of the last update or refresh (in *MMM DD, YYYY hh:mm:ss AM/PM* format) in the top-right corner of the tab. For Example: May 12, 2024 12:42:43 PM
- Update the page by clicking the refresh icon.
- On the Affected Items and Distribution tab, you can also:
 - Sort, resize, or re-arrange columns in a table (grid).
 - Show or hide columns in the table or reset page preferences, using the vertical ellipsis menu.
 - Search by using keywords—Click the search icon (magnifying glass), enter the search term in the text box, and press Enter. The search results are displayed on the same page.
 - Filter the data displayed in the table—Click the filter icon (funnel) and select whether you want to show or hide advanced filters. You can then add or remove filter criteria, save criteria as a filter, apply or clear filters, and so on. The filtered results are displayed on the same page.

Sections on the Page

The following topics describe the sections on the TV & Multicast page.

Gauge Chart

The gauge chart displays the overall health of measurements in terms of percentage. Hover over the chart to view the percentage of healthy measurements in relation to their overall health. The percentage of healthy measurements is also displayed below the chart. The percentage is 0 even if one measurement in the network is unhealthy or has an alert associated to it.

List of Plug-ins and KPIs

The plug-ins, IPTV, OTT, and Netflix Speedtest, are listed next to the gauge chart. The overall health of a plug-in is displayed next to the plug-in itself. Click any plug-in to view more KPIs that affect the overall health of that plug-in. For more information on plug-ins, see [Table 263 on page 947](#).

For example, click **Netflix Speedtest** to view other KPIs (such as Events, Metrics, and No data) that contribute to the overall percentage that is displayed. You can also view a second level and third level KPIs that affect the overall measurement health. For example, click the KPI **Metrics** to view metrics such as bandwidth speed and latency related to the plug-in you selected, which are another related KPIs that affect the overall measurement health.

A total of four plug-ins are listed by default next to the gauge chart. To clear all selected KPIs, click **Clear Selection**.

Overview Tab

The Overview tab has the Measurements widget.

The Measurements widget displays the overall percentage of healthy and unhealthy measurements in a donut chart. However, when you select a plug-in from the list next to the gauge chart, the Measurements widget displays the overall health of the measurements in relation to that plug-in. The total number of measurements in the network is displayed in the middle of the chart. Healthy measurements are highlighted in green in the donut chart, and unhealthy measurements are highlighted in red. Hover over the green portion of the donut chart to view the percentage and total number of healthy measurements. Hover over the red portion of the donut chart to view the percentage and total number of unhealthy measurements.

Timeline Tab

The Timeline tab displays the percentage of healthy measurements for a predefined (a week, a day, 3 hours, 1 hour, 30 minutes) or a custom period on a graph. By default, performance for the past 30 minutes is displayed. To change this period, click the **Week**, **Day**, **3 hrs**, **1 hr**, **30 mins**, or **Custom** buttons provided above the graph.

You can also use the range picker provided above the graph to view information for a predefined or custom period. You can click the arrows provided on either side of the range picker to move between

ranges. For example, click the **Week** button to view device health information for a week. Click the arrow on the left of the range picker to view information for the previous week. Click the arrow on the right of range picker to view information for the following week.

Data lines for the overall measurements and measurements in relation to individual plug-ins are displayed on the graph by default. Hover over any point on the line graph to view the percentage health for that particular point. When you select a plug-in from the list next to the gauge chart, the graph is updated to only show the overall health of the device in relation to that plug-in. Click any KPI to view a data line in the assigned color on the graph. Click **Clear Selection** to clear all selected KPIs.

Affected Items Tab

The Affected Items tab consists of the Measurements, Test Agents, and Devices tabs.

- The Measurements tab displays details of unhealthy measurements in the order of severity of the alerts generated. The various alerts are Action Needed and Urgent Action needed.
- The Test Agents tab displays details of Test Agents associated with the unhealthy measurements in the order of severity of the alerts generated. The various alerts are Action Needed and Urgent Action needed.
- The Devices tab displays details of devices associated with the unhealthy measurement in the order of severity of the alerts generated. The various alerts are Action Needed and Urgent Action needed.

Distribution Tab

The Distribution tab consists of the Metrics and Sites tabs.

The Metrics tab displays the list and details of unhealthy measurements grouped by metrics. The details on this tab are updated based on the KPIs that you select. You can view details such as name, affected measurements, total measurement, and health.

The Sites tab displays the details of unhealthy measurements grouped by sites. A site is the location of a device in an organization. The details on this tab are updated based on the KPIs that you select. You can view the name and the total number of affected measurements.

Troubleshoot Devices

IN THIS CHAPTER

- [Troubleshoot Using Alerts and Alarms | 450](#)
- [About the Troubleshoot Devices Page | 454](#)
- [Upgrade Images on Devices | 460](#)
- [About the *Device-Name* Page | 461](#)
- [About the Chassis Tab | 464](#)
- [About the Interfaces Tab | 466](#)
- [About the Events Page | 470](#)
- [Manage Event Templates | 480](#)

Troubleshoot Using Alerts and Alarms

The observability use case enables you (Super User, Network Admin, and Observer) to monitor the health and performance your network. Routing Director detects and generates alerts and alarms for issues in your devices and in your overall network. Timely detection of issues enables you to fix them immediately and minimize the impact of such issues on your network and its performance. We refer to alerts and alarms collectively as events in the GUI and in this topic.

You can view the events generated in your network on multiple pages in the GUI. In addition, you can also configure Routing Director to send you notifications for the events on external applications such as e-mail and Slack. If you are monitoring the network and its components in the GUI, you can drill down to the device level to view all events on the device and in its connectivity. If you received external notifications for events in the network, you can use the information in the notification message to identify the network issues. You can then determine the required fix to remediate the issue.

Navigate to the following pages in the GUI to monitor your network performance on the basis of the events generated on the devices:

- **Observability > Events > Alerts** tab

On this tab, you can view alerts related to interface, hardware, routing, and connectivity categories. You can acknowledge an alert if you have seen and taken note of the alert condition and have determined the fix for the issue raised by the alert. In addition, if you want to monitor specific alerts or alert categories, click **Templates Configuration** to create an alert template for the required alerts. You can also choose to receive and to send alert notifications over e-mail and Slack to site and organization administrators and other selected users.

Use this page during the life cycle management of your network. Observers can view the Alerts tab but cannot view or create event templates.

For more information and for a detailed list of tasks that you can perform from this tab, see ["Alerts Tab" on page 471](#).

- **Observability > Events > Alarms tab**

On this tab, you can view hardware alarms of all severities generated on the devices. You can acknowledge an alarm if you have seen and taken note of the alarm condition and have determined the fix for the issue raised by the alarm. In addition, if you want to monitor specific alarm categories, click **Templates Configuration** to create an alarm template. You can also choose to receive and to send alarm notifications over e-mail and Slack to administrators and other selected users.

Use this page during the life cycle management of the devices in your network. Observers can view the Alarms tab but cannot view or create event templates.

For more information and for a detailed list of tasks that you can perform from this tab, see ["Alarms Tab" on page 474](#).

- **Observability > Events > Device Logs tab**

On this tab, you can view the system log messages generated by the devices in your network. The device logs includes routine operations, failure or error conditions, and emergency or critical events. You can use REST APIs to search and count the logs generated per device. The tab displays logs such as total, critical, error, and warning syslogs.

Use this page during the life cycle management of the devices in your network. Observers can view the Device Logs tab but cannot view or create event templates.

For more information and for a detailed list of tasks that you can perform from this tab, see ["Device Logs Tab" on page 477](#).

- **Observability > Troubleshoot Devices page**

On this page, you can view the devices and the number of devices that have events. Alerts and alarms are displayed for all issues that require user intervention or are being monitored. You can also view a comparison of the events raised in the current week against the events of the previous week in your network. The comparison gives you an insight into network performance over short periods. This page provides you with an easy way to identify issues and drill down to the cause of the issues,

enabling you to resolve issues quickly. Use this page during the life cycle management of the devices in your network.

For more information and for a detailed list of tasks that you can perform from this page, see ["About the Troubleshoot Devices Page" on page 454](#).

- **Observability > Troubleshoot Devices > *Device-Name* > Overview** tab

On the Troubleshoot Devices page, click a device hostname to view the ***Device-Name* > Overview** tab. The accordions in the Overview tab lists the results of the tests that Routing Director runs to monitor the health of devices. Events are categorized and displayed under an accordion corresponding to that category. If a device has an issue, the severity of the event is displayed to the right of the accordion name. If there are multiple events of varying severities, the highest severity of the events is displayed.

To view more information on the events, click the > icon to the left of the accordion name to expand the accordion view. The two latest events of the highest severity are displayed on the right of the accordion under Relevant Events. If there are *fewer than two* events, hover over **View Details** for each event to view more information. If there are *more than two* events, click **View All Relevant Events**. The Events for *Device-Name* page appears and displays the complete list of events in the corresponding accordion category. You can view all the events displayed in their corresponding accordion categories and remediate the issues that may need user intervention.

Additionally, you can troubleshoot further to get more detailed information on events from the following accordions:

- **Identity & Location**—Click the compliance score of the device to view more information on the score and troubleshoot issues. For more information, see ["Identity and Location Data of a Device" on page 270](#).
- **Remote Management**—Click the **Syslog** and **Alarms** links to navigate to the **Observability > Events > Device Logs** and **Observability > Events > Alarms** pages respectively. You can view the device system logs and find more information on alarms from these pages. Additionally, if required, click **Release Router** to release the device from being managed by Routing Director. For more information, see ["Remote Management Data and Test Results" on page 272](#).
- **Hardware**—Click the data-link of an unhealthy Hardware component in the hardware accordion. The Hardware details for *Device-Name* page appears. The graphs on this page display the performance of the hardware components graphically. You can also view information on events on these performance graphs. Click the toggle button next to the name of the hardware component, to view the performance of the component in a graph. To view the details of anomalies, click the red diamond icon, orange square icon, or yellow triangle icon on the graph. The details of the anomaly appear in a pop-up. You can also zoom into a particular portion of the graph to view more information about events that have occurred.

Similarly, you can also monitor anomalies in the temperature of the chassis components from the hardware accordion. For more information on all the drill-in views available and to compare multiple graphs, see ["Hardware Data and Test Results" on page 278](#).

- **Interfaces**—Click the data-link of an unhealthy interface in the Interface accordion. The Interface details for *Device-Name* page appears. The graphs on this page display the link state and link flapping issues related to physical interfaces of the device. Click the toggle button next to the name of the interface, to view details on link state performance and issues for that interface on a performance graph. For more information on all the drill-in views available and to compare multiple graphs, see ["Interfaces Data and Test Results" on page 287](#).
- **Software**—The Software accordion enables you to fix device software version issues directly from this page. If the software on the device is out of compliance, or has reached end of life (EOL), or is approaching EOL, an alert is generated. You can fix the alerts related to the software version by clicking **Upgrade Software** to upgrade your device software to the latest recommended version. For more information, see ["Software Data and Test Results" on page 306](#).
- **Configuration**—View the most recent compliance score recorded for the device configuration. Click the score to view the compliance scan results and details about the rules that did not meet the criteria, defined in the Benchmarks document, on the **Trust > Compliance > Rule Results** page. The Benchmarks document consists of compliance policies and rules defined by Center for Internet Security (CIS). For more information, see ["Configuration Data and Test Results" on page 308](#).
- **Connectivity**—Routing Director automatically runs connectivity tests using test agents on your network devices. The results of the tests are displayed in the Connectivity accordion. Click **Retest** to re-initiate a connectivity test from this accordion. Click the data-link of an unhealthy connectivity parameter to view more information on events. The Connectivity Details page appears. Click **View all Relevant Events** to view events generated for all connections.

Additionally, you can view details of faulty connections on the topology map. Faulty connections appear as red diamond icons on connection lines. Hover over the count icon to obtain details of the faults for a connection type. The Connections table displays details about the connectivity tests run on the device. To view the test results, click the connectivity status (ERROR, PASSED, or FAILED) in the Connections table.

For more information, see ["Device Connectivity Data and Tests Results" on page 343](#).

Use this page during the life cycle management of the devices in your network. For more information and for a detailed list of tasks you can perform from this tab, see ["About the Device-Name Page" on page 461](#).

- **Intent > Device Onboarding > Put Devices into Service** page

On this page, you can view the devices and the number of devices that need user intervention to fix the issues causing the alerts and alarms. This page provides you with an easy way to identify issues

and drill down to the cause of the issues, during onboarding of devices into your network. To drill down to the issues, click a device hostname to navigate to the *Device-Name* page.

For more information and for a detailed list of tasks you can perform from this page, see ["About the Put Devices into Service Page" on page 264](#).

- **Intent > Device Onboarding > Put Devices into Service > *Device-Name*** page

On the Put Devices into Service page, click a device hostname to view the *Device-Name* page. The *Device-Name* page lists the results of the tests, that Routing Director runs to monitor the health and connectivity of a device during onboarding of the device. Device data and events are categorized and displayed under their corresponding accordions. Use this page during onboarding of devices into your network. The functionality of this page is similar to that of the **Observability > Troubleshoot Devices > *Device-Name* > Overview** tab.

For more information and for a detailed list of tasks you can perform from this tab, see ["Device Onboarding Test Results" on page 268](#).

- Device tab and Site tab in the network information table on the **Network > Devices & Links** page

View the severities of all events on the sites in the Site tab and view all the events on the devices in the Device tab. Additionally, from the Device tab, you can drill down to the device to view more information. Click a device hostname with an event on it to navigate to the **Observability > Troubleshoot Devices > *Device-Name* > Overview** tab. Use this page to view more information on the event and troubleshoot the event. Use this page during the life cycle management of the devices in your network.

For more information, see ["About the Device Tab" on page 516](#) and ["About the Site Tab" on page 524](#).

RELATED DOCUMENTATION

[Observability Overview | 393](#)

[Routing Director Overview | 3](#)

About the Troubleshoot Devices Page

IN THIS SECTION

● [Tasks You Can Perform | 455](#)

To access this page, click **Observability > Health > Troubleshoot Devices**.

Troubleshooting network issues is an important feature of the observability use case. Routing Director notifies you about significant events and anomalies within the network through alerts and alarms. You can use the information in the alert and alarm notifications to fix the anomalies and minimize the impact of the issues on the network.

The Troubleshoot Devices page provides you (superusers and network administrators) with a convenient way to monitor the health and connectivity of network devices, troubleshoot events, and manage device configurations. This page provides a summarized view of the events generated by the device and the urgency of actions required to remediate the issues causing the events.



NOTE: An observer can monitor the health and connectivity of network devices and troubleshoot events but cannot manage device configurations.

The widgets on top of the table in the Troubleshoot Devices page display the following information:

- **Urgent Action Needed**—The number of critical alerts that need urgent attention. It also displays a comparison (as a number or percentage) of alerts generated in the current week against those in the past week. Hover over the widget to view the number of critical alerts generated in the current week and in the past week.
- **Action Needed**—The number of major alerts that need attention. It also displays a comparison (as a number or percentage) of major alerts generated in the current week against the alerts in the past week. Hover over the widget to view the number of major alerts generated in the current week and in the past week. While these alerts do not require immediate attention, they do require user intervention eventually to fix the issues causing them.
- **Connected**—The number of devices in inventory that are connected to Routing Director.
- **Disconnected**—The number of devices in inventory that are not connected to Routing Director.

Tasks You Can Perform

You can perform the following tasks from this page:

- View the details of devices managed by Routing Director.

- Select the device and click **More > Detail** or hover over the device hostname and click the Details icon that appears. The Device Details pane appears on the right, displaying general information and the site information about the device. Click the close (x) icon to close the pane.
- Click the hostname of a device and the *Device-Name* page appears. The *Device-Name* page consists of the **Overview** and **Inventory** tabs.

In the Overview tab, you can view the results of the tests that Routing Director executes to determine the health and connectivity of the network devices. You can monitor device health and view the details of alerts and alarms in the accordions on this page.

In the Inventory tab, you can view the information about the hardware components of the chassis and associated interfaces, licenses applied on the device, and the features available on the licenses.

- Export inventory details of a device as a CSV file— To export a device's inventory details as a comma-separated value (CSV) file, select one or more devices and click **Export > Export as CSV**. A CSV file is downloaded to your local system.

The CSV file contains information about the hostname, IPv4 address, IPv6 address, model, serial number, OS version, type, connection status, and site of the device.

- Assign a device to a site—Sites are the physical location that host devices, such as routers, switches, and firewalls within an organization network.

To assign a device to a site:

1. Select one or more devices and click **More > Assign to Site**.

The Assign Devices to Site page appears.

2. Select the site to which you want to assign the devices from the drop-down list.
3. Click **OK**.

A message confirming that the device is assigned to the selected site appears and the site is displayed under the Site column in the Troubleshoot Devices page. The connection status of the device is Yes in the Connected column when the device is assigned to a site.



NOTE:

- You must assign a device to a site to view the statistics and inventory data of that device.

- You can perform operations like reboot, back up, open CLI, upgrade the image only on the devices that are assigned to sites.

- Reboot a device—Rebooting is the process by which a running device is restarted. You can reboot a device when there are connection or operational errors on the device.

To reboot a device:

1. Select one or more devices and click **More > Reboot**.

A reboot confirmation message appears.

2. Click **OK**.

A message indicating that the reboot has started is displayed with a link to the Audits Logs (**Settings Menu > Audit Logs**) page. You can view the progress of the reboot job in the Audit Logs page.



NOTE: You can reboot a device only if the device is assigned to a site and the connection status is Yes in the Connected column.

- Back up a device configuration—The backup operation retrieves a device's configuration and stores it in a configuration file in the database. You can use this file to restore a device's configuration in case of faulty device configurations.

To back up a device configuration:

1. Select one or more devices and click **More > Backup**.

A backup confirmation message appears.

2. Click **OK**.

A message confirming that the backup is successful appears and two links are displayed, which will redirect you to:

- The Configuration Backups (**Inventory > Devices > Configuration Backups**) page where you can view the list of backed-up device configurations. See ["About the Configurations Backups Page" on page 372](#).
- The Audit Logs (**Settings Menu > Audit Logs**) page where you can track the progress of the backup operation. See ["About the Audit Logs Page" on page 138](#).



NOTE: You can back up a device configuration only if the device is assigned to a site and the connection status is Yes in the Connected column.

On the Configuration accordion (**Observability > Health > Troubleshoot Devices > *Device-Name***), you can compare an active version of the configuration committed on a device against other backed-up versions of the same device. For more information, see ["Configuration Data and Test Results" on page 308](#).

- Upgrade the image on a device. See ["Upgrade Images on Devices" on page 460](#)

Alternatively, you can upgrade the image installed on a device by clicking the **Upgrade** button on the Software accordion (**Observability > Health > Troubleshoot Devices > *Device-Name***). For more information, see ["Software Data and Test Results" on page 306](#).

- Filter the devices—You can filter the devices based on:
 - the severity of events such as Urgent Action Needed, Action Needed, and Being Monitored, or the Healthy status of the devices.
 - the site where you deployed the devices.
- You can also perform the following tasks on this page:
 - Sort, resize, or re-arrange columns in a table (grid).
 - Show or hide columns in the table or reset page preferences, using the vertical ellipsis menu.
 - Search by using keywords—Click the search icon (magnifying glass), enter the search term in the text box, and press Enter. The search results are displayed on the same page.
 - Filter the data displayed in the table—Click the filter icon (funnel) and select whether you want to show or hide advanced filters. You can then add or remove filter criteria, save criteria as a filter, apply or clear filters, and so on. The filtered results are displayed on the same page.

For more information, see ["GUI Overview" on page 6](#).

- Set the number of rows to be displayed in the table—Set the number of rows (15, 30, 50, or 100) to be displayed in the Network Implementation Plan table in the **Display** drop-down list provided at the bottom right-corner of the table.

If there are multiple pages of rows, use the arrow heads to traverse through the pages or enter a page number to view the entries listed on that page.

Field Descriptions

[Table 113 on page 459](#) describes the fields on the Troubleshoot Devices page:

Table 113: Fields on the Troubleshoot Devices Page

Field	Description
Connected	Indicates whether the device is connected to Routing Director: • Yes—The device is connected to Routing Director. • No—The device is not connected to Routing Director.
Hostname	The hostname of the device.
Severity	Indicates the seriousness of the events on the device. The severity of the events are categorized as: Urgent Action Needed—Indicates that a critical event has occurred on the device. The functioning of the device is affected and needs urgent user intervention to fix the issue. Action Needed—Indicates that a major event has occurred on the device and needs user action but not urgently. The functioning of the device is affected but not drastically. Being Monitored—Indicates that a minor event has occurred on the device but needs no user action. The device is being monitored. Healthy—Indicates that the device is healthy without any issues.
IPv4 address	The IPv4 address assigned to the device.
IPv6 address	The IPv6 address assigned to the device.
Model	The model of the device.
Serial Number	The serial number of the device.
OS version	The OS version of the device.
Type	The type of the device.

Table 113: Fields on the Troubleshoot Devices Page *(Continued)*

Field	Description
Site	The site on which the device is deployed.

RELATED DOCUMENTATION

[About the Device-Name Page](#) | 461

Upgrade Images on Devices

SUMMARY

Use this procedure to upgrade the software image installed on one or more devices.

Device image upgrade ensures that all the devices in your network run at peak efficiency, leveraging the latest features and security enhancements. A device image can be upgraded only to a version that is uploaded to Routing Director. To upload a required software image to Routing Director, see ["Upload a Software Image" on page 366](#).



NOTE: In this release, you can upgrade images only on devices of the same model at the same time.

To upgrade image installed on one or more devices:

1. Click **Observability > Health > Troubleshoot Devices**.

The Troubleshoot Device(s) page appears.

2. Select one or more devices of the same model for which you want to upgrade the image and click **More > Upgrade**.

The Upgrade Device(s) page appears listing the devices that you selected for upgrading image.

3. Select the image to which you want to upgrade from the **Select Image** drop-down list.

The Select Image drop-down list lists the images that are relevant to the selected device model.

4. (Optional) If for some reason you don't want to upgrade the image on a device, select the device and click the **Delete** icon (dustbin) present above the list.
5. Click OK to start the upgrade process. A message confirming that the upgrade request is successful is displayed along with a link to the Audit Logs (**Settings Menu > Audit Logs**) page. You can view the progress of the image upgrade on the Audit Logs page.

You can verify that the images are upgraded from the OS Version field on the Troubleshoot Devices page.

About the *Device-Name* Page

IN THIS SECTION

- Overview Tab | 461
- Inventory Tab | 464

Use the *Device-Name* page to view the overview and inventory details of a device.

To access the *Device-Name* page:

1. Select **Observability > Health > Troubleshoot Devices**.

The Troubleshoot Devices page appears.

2. Click a device hostname.

The *Device-Name* page appears displaying the Overview and Inventory tabs.

Overview Tab

The Overview tab provides an overall view of the results of health checks that Routing Director performs on network devices. These checks are executed to assess the health of various components, such as hardware, software, interfaces, and routing. In addition, Routing Director checks the connectivity of the device and its compliance with Center for Internet Security (CIS) benchmarks. Based on the assessment of the health checks, alerts and alarms are generated for the network and listed on the **Overview** tab.

The events generated from the health checks are classified into different accordion categories. You can use these accordions to view the test results and data collected from the device and its network connectivity parameters. If a device has an issue, the severity of the issue is displayed to the right of the accordion name in the accordion corresponding to the issue. If multiple events with varying severities

occur, the highest severity level of the events is displayed. If there are no events in an accordion category, the status is displayed as Healthy.



NOTE: The severity of the events is not displayed on the Remote Management accordion.

The severity level of the events are categorized as:

- **Urgent Action Needed**—Indicates that a critical event has occurred on the device. The functioning of the device is affected and needs urgent user intervention to fix the issue.
- **Action Needed**—Indicates that a major event has occurred on the device and needs user action but not urgently. The functioning of the device is affected but not drastically.
- **Being Monitored**—Indicates that a minor event has occurred on the device but needs no user action. The device is being monitored.

You can also navigate to Juniper Support Portal by clicking the [Tech Support](#) link and to the Juniper Documentation portal by using the [View Documentation](#) link.



NOTE: For third-party devices, you cannot navigate to Juniper Support Portal and Juniper Documentation portal from the GUI.

You can click and expand each accordion to view more information about events. The alerts and alarms are displayed in the Relevant Events section on the right of the accordion. In the relevant events section, the severity of the last two events are displayed.

If there are fewer than two events, hover over **View Details** to view more information about those events. A pop-up containing information about the device name, description, creation time, the last received time, and the recurrence details of the event appears.

If there are more than two events, click **View all Relevant Events**. The Events for *Device-Name* page appears and displays the complete list of events in the corresponding accordion category. The page displays the event details such as severity, timestamp, type, recurrence, and description of each events.

The accordions on this page display the following information:

- **Identity and Location**—View general information about the device and the location where the device is installed. You can also view the most recent compliance score recorded for the device and the percentage change from the previous week's compliance score. See ["Identity and Location Data of a Device" on page 270](#).
- **Remote Management**—Displays the result of health checks made on the management connection between the device and Routing Director. You can also view details about the system log and latest alarm that the device generated, and the status of the synchronization between the device's clock

and the NTP server. A superuser can release the device from Routing Director's management from this accordion. See ["Remote Management Data and Test Results" on page 272](#).

- **Hardware**—View the temperature of the chassis and the key performance indicators (KPIs) of all the hardware components. You can view the Security Incident Resource Team (SIRT) advisories for the device that lists the vulnerabilities that affect the device. Click the data-link of an unhealthy hardware component and the Hardware details for *Device-name* page appears. On this page, you can view a graph of performance, threshold levels, events, and anomalies of the hardware components. See ["Hardware Data and Test Results" on page 278](#).
- **Interfaces**—View the power transmitted and received for optical pluggables and data related to incoming and outgoing traffic at the interfaces. Click the data-link of an unhealthy pluggable, the Pluggable for Device-Name page appears and displays the health and functioning of the pluggables. Click the data-link of an unhealthy interface, the Interface for Device-Name page appears and displays information about link state performance and issues, and port flapping issues related to physical interfaces. Click the data-links of input and output traffic to view information and graphs on input and output traffic flow through the interfaces. See ["Interfaces Data and Test Results" on page 287](#).
- **Software**—View data such as vendor, software version, device model, SIRT advisories, and so on related to the software installed on the device. The alerts are generated if the device's software is out of compliance, has reached end of life (EOL), or is nearing EOL. You can also view the reliability status of the software on the device. Superusers and network administrators can upgrade a device image from the Software accordion. See ["Software Data and Test Results" on page 306](#).
- **Configuration**—View the compliance score of the active configuration and events related to the configuration. You can view details about the rules that did not meet the criteria as defined in the Benchmarks document. Benchmarks documents consist of compliance policies and rules defined by Center for Internet Security (CIS). You can also view the overall compliance of the configuration committed on the device. Superusers and network administrators can also back up a device configuration. See ["Configuration Data and Test Results" on page 308](#).
- **Routing**—View the total number of available and unhealthy BGP, OSPF, IS-IS, RSVP, and LDP neighbors. The status on the right of the accordion displays the events related to routing and forwarding on the device. See ["Routing and MPLS Data and Test Results" on page 310](#).
- **Connectivity**—View connectivity data of all connections from the device to neighbors, Internet endpoints, cloud providers, and edge devices. Click the data-link of an unhealthy connectivity parameter to view the Connectivity Details page. On the Connectivity Details page, you can view the connections from the device in a topology view and click **View all Relevant Events** to view events generated for all connections. You can also run HTTP, DNS and ping tests for all or selected connections by clicking on **Retest**, and view the test result status and log messages for the device. To view the test results, click the connectivity status (PASSED or FAILED) in the Connections table. See ["Device Connectivity Data and Tests Results" on page 343](#).

For more information on all the accordions, see ["View Results of Automated Device Tests" on page 268](#).

Inventory Tab

Use the Inventory tab to view details about the hardware components of the chassis and associated interfaces, information about licenses applied on the device, and the features available on the licenses. To view inventory of a device, click the Inventory tab on the *Device-Name* page. The inventory tab displays the following information:

- **Chassis**—View the list of all the hardware components present on the chassis, and the associated physical interfaces. For more information, see ["About the Chassis Tab" on page 464](#).
- **Interfaces**— View details of the interfaces present on the chassis, line cards, FPCs, and PICs. You can view and modify the administration status and the description of the interfaces. For more information, see ["About the Interfaces Tab" on page 466](#).
- **Licenses**—View details about the licenses applied on the device, and information on the number of features available per license. You can add, remove, filter, and sort licenses from this page. For more information, see ["About the Licenses Tab" on page 360](#).
- **Features**—View the inventory of the features associated with the licenses that you added. Routing Director requires that you add a license to activate a feature for a device. For more information, see ["About the Features Tab" on page 358](#).



NOTE: For third-party devices, Routing Director does not support collection of inventory-related information and the Inventory tab is not displayed.

RELATED DOCUMENTATION

[About the Troubleshoot Devices Page | 454](#)

About the Chassis Tab

IN THIS SECTION

- [Tasks You Can Perform | 465](#)
- [Field Descriptions | 465](#)

To access this page from the Routing Director GUI, click **Observability > Health > Troubleshoot Devices > *Device-Name* > Inventory > Chassis**.

Use the Chassis tab to view the list of all the hardware components present on the chassis, and the associated physical interfaces.

Tasks You Can Perform

You can perform the following tasks from this page:

- View the details of hardware components present on the chassis.
- Show or hide columns in the table or reset page preferences, using the vertical ellipsis menu.

For more information, see ["GUI Overview" on page 6](#).

Field Descriptions

[Table 114 on page 465](#) describes the fields on the Chassis tab:

Table 114: Fields on the Chassis Tab

Field	Description
UUID	The UUID of the hardware component.
Module	Name of the hardware component (FPC, line card, midplane, and so on).
Model	Model of the hardware component.
Version	The software OS version of the device and the revision level of the chassis components.
Part Number	Part number of the hardware component. Built-in indicates that the hardware component is a part of the parent component and does not have a part number.
Serial number	Serial number of the hardware component. Built-in indicates that the hardware component is a part of the parent component and does not have a serial number.
Physical Interfaces	The number of physical interfaces associated with the hardware component.

Table 114: Fields on the Chassis Tab (*Continued*)

Field	Description
Description	Brief description of the hardware component.

About the Interfaces Tab

IN THIS SECTION

- [Tasks You Can Perform | 466](#)
- [Field Descriptions | 468](#)

To access this page from the Routing Director GUI, click **Observability** > **Troubleshoot Devices** > *Device-name* > **Inventory** > **Interfaces**.

Use the Interfaces tab to view details of the interfaces present on the chassis, line card, FPC, and PICs. The details include interface name, MAC address, operational status, speed, duplex mode, and so on. You can view and modify the interface name, the IP address (IPv4 or IPv6), the administration status, and the description of the interfaces.



NOTE: The interfaces information such as MAC address, operational status, speed, and duplex mode for a device is updated automatically whenever Routing Director detects a change in the interfaces configured on the device.

Tasks You Can Perform

You can perform the following tasks from this page:

- View the details of interfaces present on the chassis, line cards, FPCs, and PICs. See [Table 115 on page 468](#).
- Edit interface name, IP address, administration status, and description of an interface:

Select the interface you want to modify and click the **Edit** (pencil) icon.

- To edit the interface name, update the interface name in the Interface Name column.
- To add or modify IP address of an interface, assign or edit an IPv4 or IPv6 address of the interface in the IPv4 or IPv6 address column.

The Edit Interface confirmation page appears. Click **OK**.

- To edit the administration status, select the Up or Down caret that appears in the Administration Status column.



NOTE: To modify the administration status from the device CLI, do one of the following:

- To set the administration status as Up (enable) use `set interface interface name enable`.
- To set the administration status as Down (disable) use `set interface interface name disable`.

- To edit the description, add or modify the description field that appears in the Description column.

Click the **check mark** button that appears below the Description column to save the changes.

- Add a tag—A tag is a label in the *key:value* format that you add to an interface so that you can group the interface based on its properties or parameters. For example, you can assign a tag `type:Internet` to interfaces that allow connection to the Internet.

To add one or more tags to an interface:

1. Select an interface and click **Add Tags**.

The Manage Tags page appears.

2. In the **Tags** field, select a tag to be applied from the drop-down list that appears.

Alternatively, you can also create a new tag by clicking the **Create New Tag** link at the bottom of the page. If you are creating a new tag, ensure that:

- The key starts with a letter and is followed by letters, numbers, periods (.), hyphens (-), or underscores (_), and is between 1 and 200 characters long.
- The value includes alphanumeric characters, periods (.), hyphens (-), and underscores (_), and is between 1 and 200 characters long.

You can view all the tags applied to an interface in the Tags column.

- Sort, resize, or re-arrange columns in a table (grid).

- Show or hide columns in the table or reset page preferences, using the vertical ellipsis menu.
- Filter the data displayed in the table—Click the filter icon (funnel) and select whether you want to show or hide advanced filters. You can then add or remove filter criteria, save criteria as a filter, apply or clear filters, and so on. The filtered results are displayed on the same page.

Field Descriptions

Table 115 on page 468 describes the fields on the Interfaces tab:

Table 115: Fields on the Interfaces Tab

Field	Description
ID	ID of the interface.
Interface Name	Name of the interface. Click the arrow that appears next to the interface name to view the logical interfaces associated with the physical interface.
IPv4 Address	IPv4 address assigned to the interface. You can configure the IP address from the device CLI by using the <code>set interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>IPv4-address/prefix-length</i></code> command.
IPv6 Address	IPv6 address assigned to the interface. You can configure the IP address from the device CLI by using the <code>set interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>IPv6-address/prefix-length</i></code> command.
MAC Address	MAC address of the interface.

Table 115: Fields on the Interfaces Tab *(Continued)*

Field	Description
Admin Status	Administration status of the interface. Super User or Network Admin can enable or disable the administration status of an interface by using the device CLI and also from the GUI. For example, you can force an interface failure by administratively disabling the interface. The values are: • Up—Interface is enabled. • Down—Interface is disabled.
Operational Status	Operational status of the interface. The values are: • Up—Physical link of the interface is operational. • Down—Physical link of the interface is not operational.
Speed	Interface speed in Mbps or Gbps. Speed is displayed only when the administration status and operational status are up.
Duplex Mode	Indicates if the duplex mode on the interface is full-duplex or half-duplex.
Description	Brief description of the interface. You can edit the description field.
Tags	Tags assigned to the interface. If there are more than one tag assigned to the interface, an integer is present next to the displayed tag. For example +2 in the Tags column indicates that there are two more tags in addition to what is currently listed under Tags. You can hover over the integer to view the additional tags assigned to the interface.

About the Events Page

SUMMARY

Users with the Super User, Network Admin, or Observer roles can use the **Events** page. The users can monitor the health of the network using notifications such as alerts, alarms, and device system logs from this page.

IN THIS SECTION

- [Alerts Tab | 471](#)
- [Alarms Tab | 474](#)
- [Device Logs Tab | 477](#)

To access this page, click **Observability > Health > Events**.

Routing Director generates notifications based on data collected from the network and network devices. These notifications highlight issues that may need attention and how they can affect the network.

Routing Director monitors Key Performance Indicators (KPIs) related to a device's health and the network connectivity parameters. When anomalies occur in the KPIs, Routing Director generates alerts to notify you of these anomalies. For example, interface input errors generate an alert.

Alarms are standard trigger conditions set for devices. They are events that indicate conditions on a device that might prevent the device from operating normally. For example, gateway device fault triggers an alarm.

The Events page has three tabs to display alerts, alarms, and device system logs. You can view and manage notifications for alerts and alarms, and view device system logs from their respective tabs. The Alerts tab is displayed by default. Alerts and alarms are collectively called events in the Routing Director GUI and in this topic. By default, the tables on the Alerts and Alarms tabs display the events based on the time they were received, with the latest event on top. On each tab, you can see three widgets that display important network and device statistics such as the total number of events generated and the number of critical and noncritical events that have recently been detected in your network.

In addition, you can view specific alerts and alarms by applying an event template to your organization. Event templates filter the alert list to display only the alerts tracked in the template. You cannot filter individual alerts by using an event template. You can only configure Webhook or e-mail notifications for individual alerts.

Note that an e-mail notification is sent out only when you have configured SMTP in Routing Director. For information about configuring SMTP, see "[Configure SMTP Settings in Paragon Shell](#)" on page 101. The default behavior is to not send e-mail notifications for events.

To send event notifications to Slack channels, configure Webhooks on the Organization Settings page (**Settings Menu > System Settings > Webhooks**). For more information, see Parameters to Configure Webhooks section of the "[Manage Organization Settings](#)" on page 60 topic.

The **Events** page displays device notifications in the following tabs:

Alerts Tab

IN THIS SECTION

- [Tasks You Can Perform | 472](#)
- [Field Descriptions | 473](#)

To access this tab, click **Observability > Health > Events**. The Alerts tab is displayed by default.

Routing Director generates various alerts to notify you of anomalies in the KPIs in your network. This tab displays all the generated alerts, by default. To monitor specific alerts, you can apply an event template to your organization. Event templates filter the alert list to display only the alerts that are tracked in the template. You can also choose to receive e-mail and Slack (using Webhooks) notifications for the alerts. For more information, see "[Manage Event Templates](#)" on page 480.



NOTE: The page auto-refreshes every one minute.

You can view the following statistics in the widgets on the Alerts tab:

- **Total Alerts**—Displays the total number of alerts generated in the organization.
- **Critical Alerts**—Displays the number of active critical alerts that need immediate attention. Examples of critical alerts include, OSPF send module is not functioning, flaps are increasing continuously, and FPC heap memory utilization exceeds the critical threshold.
- **Major Alerts**—Displays the number of active major alerts generated in the organization. While these alerts do not require immediate attention, you will eventually need to fix the issues that cause them. Examples of major alerts include a low trust score.
- **Minor Alerts**—Displays the number of active minor alerts generated in the organization. They are warnings that needs to be fixed but don't require immediate attention. Examples of minor alerts include, system power remaining is 50 percent and temperature has exceeded default warning threshold.



NOTE: Active alerts are alerts that currently exist on the device and are not yet acknowledged or fixed. The status of active alerts is shown as Open.

Tasks You Can Perform

- Create event templates—Click **Templates Configuration** to create one or more event templates. For more information, see ["Create an Event Template" on page 481](#).
- View details of an alert—Select an alert and click **More > Detail** or click the **Details** icon on the left to view more information on the alert. The **Alert Details** pane appears displaying the alert ID, alert group, acknowledge or unacknowledge time, and acknowledgment note.



NOTE: You can drill-down to the device level to view more details on the alert. Click a Device name next to an alert to navigate to the Overview tab of the **Troubleshoot Devices > Device-Name** page. On the Overview tab one of the following health status is displayed (on the right) for each accordion:

- Healthy
- Urgent Action Needed (Critical)
- Action Needed (Major)
- Being Monitored (Minor)

You can click the accordions and analyze the issues that have occurred on the device. The Relevant Events section provides additional insights on the events.

- Acknowledge an alert—When you want to mark an issue raised by one or more alerts as seen, you can mark it as acknowledged.

You can acknowledge an alert to indicate that the issue raised has come to your notice. Acknowledging an alert doesn't mean that the issue is fixed. For example, during a maintenance window, multiple alerts are raised. But not all of them will prevent the devices from operating normally. In such cases, you can acknowledge those minor alerts but you won't necessarily have to take any corrective actions.

To acknowledge alerts, select one or more alerts and click **More > Ack**. The **Acknowledge** confirmation window appears. Enter an acknowledgment message in the **Note** field and click **OK**. Once acknowledged, the status of the alert is changed to **Ack**.



NOTE: After you have acknowledged an alert, there might be a delay (of a few seconds) for the status of the alert to be updated.

- Unacknowledge an alert—If you acknowledged an alert by mistake and want to reverse that operation, you can unacknowledge the alert.

To unacknowledge alerts, select one or more alerts and click **More > Unack**. The **Unacknowledge** confirmation page appears. Once unacknowledged, the status of the alert is changed to **Open** and the note will now be cleared.

- Hide acknowledged alerts—Select the **Hide Acknowledged** check box to hide the acknowledged alerts in the alerts table. The table is then updated to display only open alerts.
- Filter the data displayed in the table—Click the filter icon (funnel) and select whether you want to show or hide advanced filters. You can then add or remove filter criteria, save criteria as a filter, apply or clear filters, and so on. The filtered results are displayed on the same page.
- Show or hide columns in the table or reset page preferences, using the vertical ellipsis menu.
- Sort, resize, or re-arrange columns in a table (grid).

Field Descriptions

[Table 116 on page 473](#) describes the fields in the Alerts Tab.

Table 116: Fields in the Alerts Tab

Field	Description
Alert ID	The unique identifier of the alert.
Type	The category of the alert.
Description	The description of the issue.
Severity	The severity level of the issue that raised the alert. Options are: • Critical—Indicates that the issue needs immediate attention. • Major—Indicates that the issue is not critical and does not require immediate attention, but will need to be fixed eventually. • Minor—Indicates that the issue is being monitored and currently there is no impact on the functioning of the network or network devices.
Raise Time	The date and time at which the alert was raised.
Site	The site in which the device (for which the alert was raised) is located.

Table 116: Fields in the Alerts Tab *(Continued)*

Field	Description
Device	The name of the device on which the alert was raised.
Services	The service instance associated with the alert.
Ack Time	The date and time at which the alert was acknowledged.
Ack Note	The reason for which the alert was acknowledged.

Alarms Tab

IN THIS SECTION

- [Tasks You Can Perform | 475](#)
- [Field Descriptions | 476](#)

To access this tab, click **Observability > Health > Events > Alarms**.

Alarms are generated by devices when an abnormal event prevents the device from functioning normally. Alarms provide information and help you monitor the status and the health of your network devices. The Alarms tab displays all the generated alarms, by default. To monitor specific alarms, you can apply an event template to your organization. Event templates filter the alarm list to display only the alarms that are tracked in the template. You can also choose to receive e-mail and Slack (using webhooks) notifications for the alarms. For more information, see "[Manage Event Templates](#)" on page 480.



NOTE: The tab auto-refreshes and displays the latest alarms.

You can view the following statistics in the widgets on the Alarms tab:

- **Total Active Alarms**—Displays the total number of alarms raised by devices in the organization. You can also view the total number of new alarms generated in the past 24 hours and in the past week. This number can vary based on the filters selected and the event template applied.

- **Critical Active Alarms**—Displays the number of critical alarms that need immediate attention. An example of a critical alarm is input voltage failure. You can also view the number of critical alarms generated in the past 24 hours and in the past week.
- **Warning Active Alarms**—Displays the number of minor alarms raised. Examples of warning alarms include minimum supported firmware version mismatches or when the host active disk usage exceeds the threshold. You can also view the number of new warning alarms generated in the past 24 hours and in the past week.

In addition to critical and warning alarms, you can also view informational alarms in the alarms table. To view informational alarms, click the filter (**funnel**) icon. From the **Field** list, select **Severity** and from the **Value** list, select **Info**. Click **Save** and **Close**. The alarms table is updated to show only informational alarms.

You can click the widgets on the page to filter the displayed alarms and alarms statistics. For example, if you click **Critical active alarms**, then the **Total active alarms** widget and the alarms table update to display the number and details of only the critical active alarms.

Tasks You Can Perform

- **Create event templates**—Click **Templates Configuration** to create one or more event templates. For more information, see ["Create an Event Template" on page 481](#).
- **View details of an alarm**—Select an alarm and click **More > Detail** or click the **Details** icon on the left to view more information about the alarm. The Alarm Details page appears displaying the alarm ID, alarm group, cleared time, acknowledge or unacknowledge time, and acknowledgment note.



NOTE: You can drill-down to the device level to view more details on the alarms. Click a Device name next to an alarm to navigate to the Overview tab of the **Troubleshoot Devices > Device-Name** page. On the Overview tab one of the following health status is displayed (on the right) for each accordion:

- Healthy
- Urgent Action Needed (Critical)
- Action Needed (Major)
- Being Monitored (Minor)

You can click the accordions and analyze the issues that have occurred on the device. The Relevant Events section provides additional insights on the events.

- **Acknowledge an alarm**—When you want to mark an issue raised by one or more alarms as seen, you can mark it as acknowledged.

You can acknowledge an alarm to indicate that the issue raised has come to your notice. Acknowledging an alarm doesn't mean that the issue is fixed. For example, during a maintenance window, multiple alarms are raised. But not all of them will prevent the devices from operating normally. In such cases, you can acknowledge those informational alarms but you won't necessarily have to take any corrective actions. You can acknowledge only open alarms.

To acknowledge alarms, select one or more alarms and click **More > Ack**. The **Acknowledge** confirmation window appears. Enter an acknowledgment message in the **Note** field and click **OK**.



NOTE: The status of the alarm remains Open and does not change when you acknowledge an alarm.

- Unacknowledge an alarm—If you acknowledged an alarm by mistake and want to reverse that operation, you can unacknowledge the alarm.

To unacknowledge alarms, select one or more alarms and click **More > Unack**. The **Unacknowledge** confirmation page appears. Enter an unacknowledgment message in the **Note** field and click **OK**. Once unacknowledged, the status of the alarm is changed to Open.

If you do not add a note and there was a previously added note for the alarm, the note will now be cleared.



NOTE: The status of the alarm remains Open and does not change when you unacknowledge an alarm.

- Hide acknowledged alarms—Select the **Hide Acknowledged** check box to hide acknowledged alarms in the alarms table. The table is then updated to display only open alarms that have not been acknowledged.
- Filter the data displayed in the table—Click the filter icon (funnel) and select whether you want to show or hide advanced filters. You can then add or remove filter criteria, save criteria as a filter, apply or clear filters, and so on. The filtered results are displayed on the same page.
- Show or hide columns in the table or reset page preferences, using the vertical ellipsis menu.
- Sort, resize, or re-arrange columns in a table (grid).

Field Descriptions

Table 117 on page 477 describes the fields in the Alarms tab.

Table 117: Fields in the Alarms Tab

Field	Description
Alarm ID	Unique identifier of the alarm.
Device	Name of the device on which the alarm occurred. You can click the <i>Device Name</i> to see in-depth device information on the Observability > Troubleshoot Devices > Device-Name > Overview tab.
Details	Details of the issue. For example, operational status of an interface is down. In most cases, the component affected by the alarm is displayed.
Severity	Severity level or seriousness of the alarm. Options are: • Critical—Indicates that the issue needs immediate attention. • Warning—Indicates that the issue needs to be fixed but doesn't require immediate attention. • Information—Indicates that the issue is being monitored but currently there is no impact on the functioning of the device.
Raised	Date and time when the alarm was raised.
Type	Category of the alarm. The alarm category is Hardware.
Site	Site in which the device (for which the alarm was raised) is located.
Status	Status of the issue that raised the alarm. Options are: • Open: Alarm is still active. • Cleared: Alarm is not active as it is fixed or closed by the device.

Device Logs Tab

IN THIS SECTION

- [Tasks You Can Perform | 478](#)
- [Field Descriptions | 479](#)

To access this tab, click **Observability > Health > Events > Device Logs**.

Devices generate system log messages to record events such as:

- Routine operations such as creation of an Open Shortest Path First (OSPF) protocol adjacency or a user login to the configuration database.
- Failure or error conditions such as failure in accessing a configuration file or an unexpected closure of a connection to a peer process.
- Emergency or critical conditions such as a router powering down due to excessive temperature.



NOTE: You can use REST APIs to search and count the logs generated per device.

The **Device Logs** tab displays all the system logs generated from the devices in your network.



NOTE: The page auto-refreshes every one minute.

You can view the following statistics in the widgets on the Device Logs tab:

- **Total Syslogs**—Displays the total number of system logs generated for all devices in the organization. This number can vary based on the filters selected.
- **Critical Syslogs**—Displays the number of critical system logs.
- **Error Syslogs**—Displays the number of error system logs.

In addition to critical and error system logs, you can also view warning system logs from the device logs table. To view warning system logs, click the filter (**funnel**) icon. From the **Field** list, select **Severity** and from the **Value** list, select **Warn**. Click **Save** and **Close**. The device logs table is updated to show only warning system logs.

You can click the widgets on the page to filter the displayed system logs and system logs statistics. For example, if you click **Critical Syslogs**, then the **Total Syslogs** widget and the device logs table update to display the number and details of only the critical system logs.

Tasks You Can Perform

- View the system logs for all devices in the organization—Select one of the following time intervals for which you can want to view the system logs:
 - Week
 - Day
 - 3hrs

- 1hr
- 30 minutes
- Custom—When you select this option, the calendar is enabled on the left. Click the calendar icon to manually select the date and time range for the past month. The logs are immediately displayed in a table.



NOTE:

- By default, logs generated in the past 30 minutes are displayed.
- System logs are collected from the device every three minutes and stored securely. The retention period for system logs is one month.

- Filter the data displayed in the table—Click the filter icon (funnel) and select whether you want to show or hide advanced filters. You can then add or remove filter criteria, save criteria as a filter, apply or clear filters, and so on. The filtered results are displayed on the same page.
- Show or hide columns in the table or reset page preferences, using the vertical ellipsis menu.
- Sort, resize, or re-arrange columns in a table (grid).

Field Descriptions

Table 118 on page 479 describes the fields in the Device Logs tab.

Table 118: Fields in the Device Logs Tab

Fields	Description
Device	The name of the device that generated the log.
Hostname	The name that identifies the device in the network. NOTE: To view the host name, hover over Show/Hide Columns and enable Hostname check box. The host name is displayed in the device logs table.

Table 118: Fields in the Device Logs Tab *(Continued)*

Fields	Description
Severity	The severity level of the event that generated the log. Options are: • Critical—Indicates that the issue needs immediate attention. • Error—Indicates that the issue needs to be fixed but doesn't require immediate attention. • Warning—Indicates that the event is being monitored but currently there is no impact on the functioning of the device.
Timestamp	The date and time at which the logged event was recorded.
Site	The site in which the device is located.
Appname	The application on the device that generated the log message.
Message	The details of the log. NOTE: If the entire message is not fully visible, you can hover over the displayed message to view the complete log message.
Raw Message	The original device log messages in unprocessed format. To view the raw device log message, hover over Show/Hide Columns and enable Raw Message check box. The raw message is displayed in the device logs table.

Manage Event Templates

SUMMARY

Users with the Super User and Network Admin roles can use the Event Templates Configuration page to create, edit, clone, or delete event templates.

IN THIS SECTION

- [Create an Event Template | 481](#)
- [Edit Event Template Configuration | 485](#)
- [Clone an Event Template | 486](#)

Routing Director allows you to create event templates to monitor and notify users about specific alerts or alarms through third party applications such as e-mail and Slack. To send notifications of the events to Slack channels, configure webhooks on the Organization Settings page (**Settings Menu > System Settings > Webhooks**). For more information, see "[Configure Webhooks to Receive Event Notifications](#)" on page 81.



NOTE: An observer can monitor events generated in the organization from the **Events** page but cannot view the **Event Templates Configuration** page or manage event templates.

The Alerts tab (**Observability > Health > Events > Alerts**) displays all active alerts generated by the system even if the alert notification has not been enabled in an event template. The conditions to raise alerts is configured separately for each use case. The Alarms tab (**Observability > Health > Events > Alarms**) only displays alarms that are enabled in an event template.

You can create one or more event templates for an organization, however, only one event template can be active at a time. In the template, specify the events that you want to track. You can also enter specific e-mail recipients who will be notified about the events detected in your network.



NOTE: You can apply only one template to an organization at a time.

Create an Event Template

To create a new event template for alerts or alarms:

1. Click **Observability > Health > Events**.

The Events page appears displaying the Alerts tab, by default.

2. Select the **Alerts** or **Alarms** tab.

If an event template is already configured and applied to the organization, a list of filtered alerts and alarms is displayed. Otherwise, all generated alerts and alarms are displayed in the respective tabs.

3. Click **Templates Configuration**.

The Event Templates Configuration page appears displaying the existing event templates list on the left. If an event template is currently applied to the organization, it is displayed at the top with a green ✓ icon (marked as *Entire Org*).

4. To create a new template, click **Create Template** and configure the fields as described in [Table 119 on page 482](#).



NOTE: Fields marked with an asterisk (*) are mandatory.



NOTE: By default, the **Event Templates Configuration** page displays details of the last template that was configured. You must click **Create Template** to add a new event template, otherwise you will be editing the last template that was configured.

5. Click **Save** to save the template.

A confirmation message appears stating that the template has been successfully created. The newly created template is listed on the left.

Alternatively, you can click **Cancel** to discard your changes.

Table 119: Fields on the Event Template Configuration Page

Field	Description
Name*	Enter a name for the new template.
Apply to Scope	Select the scope of this template (if you want to apply to the organization or not). If you enable the Mark as Active check box, this template configuration will be active and applied to the organization. Activating the template enables the selected events and sends e-mail and webhook notifications for only those events selected in the template. NOTE: You can apply only one template to an organization at a time.

Table 119: Fields on the Event Template Configuration Page *(Continued)*

Field	Description
Email Recipients Settings	Enable the respective check box to send event notification e-mails to all organization administrators or all site administrators. All administrators must enable e-mail notification settings in their accounts to receive event notification e-mails. - To enable e-mail notifications at the organization-level, click My Account > Email Notifications > Enable > Enable Org Notifications. - To enable e-mail notifications at the site-level, My Account link > Email Notifications > Enable > Enable Org Notifications. Enable the toggle button next to a site to receive e-mail notifications specific to that site. NOTE: Administrators who do not enable this setting in their accounts will not receive any event notifications. For more information, see Enable E-mail Notifications.
Additional email recipients	Enter the e-mail addresses of recipients who will receive the event notification e-mails. The recipients can be users of the organization and outside the organization.

Table 119: Fields on the Event Template Configuration Page *(Continued)*

Field	Description
Event Types*	Click Alerts to view alert event types. • Select one or more check boxes under Send Webhook Notification to receive a notification when the selected alert is generated. • Select one or more check boxes under Send Email Notification to send e-mail notifications when the selected alerts are generated. You can select one or more individual alert types, or select a category to select all alert types under the category. Alert categories are Interface, Hardware, Routing, and Connectivity. Click Alarms to view alarm event types. • Select one or more check boxes under Enable Alarm to monitor the alarms listed in this template. • Once you select the alarm types, you can also enable the corresponding check boxes under Send Email Notification to send e-mail notifications when the selected alarms are generated. You can select one or more individual alarm types, or select the category to select all alarm types under the category. Alarm category is Hardware. Options for alarm types are: • Gateway Device Fault (alarm raised on a Juniper gateway device.) • Switch Device Fault (alarm raised on a Juniper switch device.) • Chassis alarms—Predefined alarms triggered by a physical condition on the device. Options are: • Switch PoE Alarm • Switch PEM Alarm • Switch Power Supply Alarm • Switch Storage Partition Alarm

Table 119: Fields on the Event Template Configuration Page (*Continued*)

Field	Description
	• Switch Fan Alarm Click Custom Alerts to view custom alert event types. • Select the Observability Custom Alert check box under Send Webhook Notification to receive a notification when the selected (custom) alert is generated. • Select the Observability Custom Alert check box under Send Email Notification to send e-mail notifications when the selected alert is generated. NOTE: The check boxes under Send Email Notification for Alerts, Alarms, and Custom Alerts are enabled only when you configure SMTP. For more information, see Enable E-mail Notifications.

Edit Event Template Configuration

To edit the configuration of an existing event template:

1. Click **Observability > Health > Events**.

The Events page appears displaying the Alerts tab, by default.

2. Click **Templates Configuration** on the top-right corner.

The Event Templates Configuration page appears displaying all the existing event templates on the left.

3. Select the template that you want to edit.

The template details are displayed.



NOTE: You can search for an existing event template from the **Search Template** field on the top left corner.

4. Edit the fields as described in [Table 119 on page 482](#).

5. Click **Save** to save the changes to the template.

A confirmation message appears stating that the template has been successfully updated. The updated template is listed on the left.

Alternatively, you can click **Cancel** to discard your changes.

Clone an Event Template

You can clone an existing template if you want to quickly create a new template by making minor changes to an existing template configuration.

To clone an existing event template:

1. Click **Observability > Health > Events**.

The Events page appears displaying the Alerts tab, by default.

2. Click **Templates Configuration** on the top-right corner.

The Event Templates Configuration page appears displaying all the existing event templates.

3. Select the template that you want to clone.

The template details are displayed.



NOTE: You can search for an existing event template from the **Search Template** field on the top left corner.

4. Click the copy icon (on the right) to clone the current template.

A notification appears stating that *Name_clone_template* is created. The cloned template is listed under the current template on the left and the details are displayed.

5. Update the fields of the cloned template as described in [Table 119 on page 482](#).

6. Click **Save** to save the changes to the template.

A confirmation message appears stating that the template has been successfully created. The cloned template is listed on the left with the configuration you specified.

Alternatively, you can click **Cancel** to discard your changes.

Delete an Event Template

To delete an existing event template:

1. Click **Observability > Health > Events**.

The Events page appears displaying the Alerts tab, by default.

2. Click **Templates Configuration** on the top-right corner.

The Event Templates Configuration page appears displaying all the existing event templates on the left.



NOTE: You can search for an existing event template from the **Search Template** field on the top left corner.

3. Select the template that you want to delete.

The template details are displayed.



NOTE: You cannot delete a template that is currently applied to the organization. To delete a template that is assigned to the organization, you must first unassign it by disabling the **Apply To Scope** field (see [Table on page 482](#)). Once the template is unassigned or if another template has been assigned to that organization, the delete option will be enabled for the template you want to delete.

4. Click the delete (trash can) icon.

A confirmation message appears.

5. Click **Yes** to delete the template.

A notification message appears stating that the template is successfully deleted. The template is immediately removed from the list on the left.

Alternatively, you can click **No** to discard your current changes.

SEE ALSO

[About the Events Page | 470](#)

Monitor Devices Using AI/ML

IN THIS CHAPTER

- [Automatically Monitor Device and Interface Health and Detect Anomalies | 488](#)
- [Detect Blackholes | 495](#)

Automatically Monitor Device and Interface Health and Detect Anomalies

SUMMARY

Use this topic to understand how Routing Director automatically monitors device health and detects anomalies, and how you can use the GUI to view anomalies related to device health.

IN THIS SECTION

- [Device and Interface Health Monitoring and Anomaly Detection Overview | 488](#)
- [View Device and Interface Health Anomalies on the GUI | 491](#)

Device and Interface Health Monitoring and Anomaly Detection Overview

IN THIS SECTION

- [RCA of Temperature Anomalies | 490](#)
- [Device Health KPIs | 490](#)



NOTE:

- Device Health and anomaly detection is a beta feature in this release.
- To monitor device health, you must enable AI/ML (`install-aiml`) and device health monitoring (`enable-device-health`) when you install the Routing Director cluster. For more information, see [Deploy the Cluster](#).

Enabling AI/ML requires additional system resources (CPU and memory). For information about the additional resources required for AI/ML, see *Hardware Requirements*.

To ascertain the health of a network, you need to monitor the health of the devices and their interfaces in the network. Routing Director uses AI/ML (artificial intelligence [AI] and machine learning [ML]) techniques to automatically monitor Key Performance Indicators (KPIs) related to a device's health, and automatically detects any anomalies that occur. Routing Director also performs a root-cause analysis (RCA) of device temperature anomalies when the device is in operation.

The periodic monitoring of the device's health status and the timely detection of device and interface health anomalies enables operators to take action and minimize the impact of any issues that occur

Routing Director monitors device health in the following scenarios:

- **During device onboarding**—When a device is being onboarded, Routing Director monitors the device's health and generates an alert if any anomalies occur.

When a device is being onboarded, if other devices of the same model that were previously onboarded exist, Routing Director compares the data to detect anomalies. However, if a device of a particular model is being onboarded for the first time, then the efficacy of the anomaly detection is limited because of lack of historical data.

- **During device operation**—After the device is onboarded successfully and is managed, Routing Director continuously monitors the KPIs related to device health. For each KPI of each device, Routing Director monitors the KPI, forecasts the range, and detects any anomalies that occur. During device operation, Routing Director detects device health anomalies (within 30 minutes) based on historical data for that device and the forecasted range.



NOTE:

- In the validation phase, the MAPE score for the ML models used in device health monitoring was observed as varying between 2.5 to 6.5.
- After a KPI value changes, the forecasted range takes approximately two hours to stabilize.

RCA of Temperature Anomalies

When a device is in operation, Routing Director provides RCA for issues related to the Routing Engine temperature and Routing Engine CPU temperature. Routing Director analyzes the different attributes (CPU utilization percentage, fan RPM percentage, and inlet temperature) that could cause a temperature issue. Routing Director also compares the device's temperature to an expected range. Based on the analysis and comparison, Routing Director provides an alert, an expected reason for the issue, and details on the events that might have caused the issue. [Figure 39 on page 490](#) displays a sample page showing the RCA logs for an anomaly in the Routing Engine temperature.

Figure 39: Sample Page Showing RCA for Device Temperature Anomaly



Device Health KPIs

[Table 120 on page 490](#) displays the device health KPIs that Routing Director monitors for each device.

Table 120: KPIs Related to Device Health

KPI	Component	Parameters
CPU	Routing Engine	CPU Utilization Percentage (%)
	Line card	

Table 120: KPIs Related to Device Health *(Continued)*

KPI	Component	Parameters
Memory	Routing Engine Line card	Memory Utilization Percentage (%)
Fan	Not applicable	RPM Percentage (%)
Temperature	- Routing Engine (RE) - Routing Engine CPU - Line card - Line card CPU	Current temperature

Table 121 on page 491 lists the KPIs related to interface health that Routing Director monitors for each interface.

Table 121: KPIs Related to Interface Health

KPI	Description
Optics Rx power Optics Tx power	Current optics power level in dBm.
Input traffic Output traffic	Current traffic in Mbps.
Optical/Module temperature	Current optics temperature in °C.

View Device and Interface Health Anomalies on the GUI

You can view and monitor the device health anomalies for a device on the Hardware accordion of the *Device-Name* page.

To view and monitor device health anomalies:

1. Do one of the following.

- To view and monitor device health and interface health anomalies during device onboarding, select **Inventory > Device Onboarding > Onboarding Dashboard > Put Devices into Service > *Device-Name***.
- To view and monitor device health and interface health anomalies during device operation, select **Observability > Troubleshoot Devices > *Device-Name***.

The *Device-Name* page appears.

2. To View:

- Device health, scroll to the Hardware accordion and click > to expand the accordion.
 - The Chassis section of the accordion displays the health status of the following KPIs monitored by Routing Director:
 - Fans
 - CPU
 - Linecards
 - Memory
 - Temperature
 - Device events appear under Relevant Events with the following information:
 - Event notification message
 - Date and time that the last event was received by Routing Director.
- Interface health, scroll to the Interface accordion and click > to expand the accordion. You can view the below KPIs:
 - Optical temperature, optical Tx power, and optical Rx power of pluggables
 - Input traffic
 - Output traffic
 - Interface events appear under Relevant Events with the following information:
 - Event notification message
 - Date and time that the last event was received by Routing Director.

3. Hover over or click **View Details** to view the details of the event, including the number of times that the event occurred.

4. (Optional) Click **View All Relevant Events** to view all the health-related events for the device.

The events appear on the Events for *Device-Name* page.

5. You can view detailed information about each KPI related to device or interface health by doing the following:

- a. Click the health status link for the KPI; for example, Fans or Temperature.

The Hardware details for *Device-Name* page appears, displaying the section for the KPI that you clicked in the preceding page.

For example, if you click the link for Fans, then the Fans section is expanded and the graphs related to the fans are displayed.

[Figure 40 on page 493](#) shows a sample section (Temperature) of the Hardware Details for *Device-Name* page.

For interfaces, click the status link of the KPI; for example, Input Traffic or Output Traffic. Respective page appears displaying the graph related to the KPI. For example, clicking Input Traffic link opens the Input Traffic details for *Device-Name* page displaying graphs related to input traffic and details of alerts if any. [Figure on page 494](#) shows a sample graph for input traffic.

See ["Interfaces Data and Test Results" on page 287](#) for details on the graphs and KPIs related to interfaces.

- b. To view the details of an anomaly, click the yellow triangle icon on the graph.

The details of the anomaly appear in a pop-up, as shown in [Figure 40 on page 493](#) for hardware and as shown in [Figure on page 494](#) for interfaces.

6. Click **Close** or the **X** icon to go to the *Device-Name* page.

For more information on the hardware accordion, see ["Hardware Data and Test Results" on page 278](#) and on interface accordion, see ["Interfaces Data and Test Results" on page 287](#).

Figure 40: Sample Hardware Details for *Device-Name* Page



1– KPI	5– High threshold marker
2– Circle icons indicating that the KPI is normal	6– Pop-up showing details of device health anomaly.
3– Upper and lower boundaries (dynamic thresholds) for the data displayed in the graph	7– Triangle icons indicating an anomaly when the higher threshold is breached.
4– Critical threshold marker	8– Legend showing the colors for different sub-components used in the graphs

Figure on page 494 shows the input traffic through et-0/0/1 interface on a device during a 30 minute interval. A warning (indicated by the yellow triangle icon) is raised to indicate an anomaly in the projected input traffic rate,

Figure 41: Input Traffic Page Showing Anomalies in measured Input Traffic



A KPI value is considered:

- Anomalous if the KPI value is outside the dynamic threshold (shaded area of the map) for nine consecutive intervals or nine minutes of data collection.
- Normal if the KPI value falls within the dynamic threshold for three consecutive intervals of data collection.

If the KPI value continues to be outside the dynamic threshold for more than nine consecutive intervals, the dynamic threshold adapts to the new values and a new dynamic threshold is created. An alert is raised if the KPI value crosses the High or Critical values irrespective of whether the value falls within the dynamic threshold or not.

Detect Blackholes

SUMMARY

Use this topic to understand how Routing Director provides notifications about blackholes and information for root cause analysis of the blackhole events.

IN THIS SECTION

- [View Blackhole Details | 496](#)



NOTE: Blackhole detection is a beta feature in this release.

A traffic blackhole occurs when transit traffic is consistently or intermittently dropped. Blackholes are caused due to:

- Hardware errors (interface, fabric, queuing drop counters)
- Software issues
- Side effects of events such as BFD Session Flaps, Link flaps, Line card restart, GRES, NSR, Config Change
- Resource exhaustion
- Incorrect configuration

To detect blackholes, blackhole detection must be enabled in Routing Director.

Use the following commands to configure blackhole detection in Routing Director:

```
set paragon cluster applications aiops install-aiml true
set paragon cluster applications aiops enable-blackhole true
```

When you enable blackhole detection on devices, playbook rules for collecting KPIs (packet drops, software errors, and hardware interrupts) to detect blackholes are applied to the devices.



NOTE: Enabling AI/ML requires additional system resources (CPU and memory). For information about the additional resources required for AI/ML, see *Hardware Requirements*.

Routing Director uses AI/ML to detect blackholes on devices onboarded to Paragon Automation. When blackhole detection is enabled on Routing Director, the sensors are initialized to capture input packet

rates, and output packet rates, and packet drop rates. When captured rates are compared to respective thresholds and blackhole is detected if the packet drop rate threshold is persistently breached.

View Blackhole Details

Routing Director identifies blackholes based on changes in the number of packets dropped over a period of time. If the number of packets dropped exceeds a specified threshold, Routing Director generates a critical alarm indicating blackhole. This information is displayed as Traffic Loss in the Alerts table (**Observability > Events > Alerts tab**) and the Routing and MPLS accordion of the *Device-Name* page. To access the *Device-Name* page:

- For a device in service, navigate to **Observability > Troubleshooting Devices** and click a device.
- For a device during onboarding, navigate to **Inventory > Device Onboarding > Onboarding Dashboard** and click a device.

Figure 42 on page 496 displays the **Blackhole Detected** alert displayed on the Routing and MPLS accordion.

Figure 42: Blackhole Detected Alert



Click the **Alerts** link of the Traffic Loss field to view graphs of packet input rate, packet output rate, packet drop rate, and traffic loss alerts. Figure on page 496 shows the Traffic Loss page with alerts indicated as follows:

- A red alert icon (critical) indicates that blackhole is detected
- An orange alert icon (major) indicates that packet drop anomalies are present but blackhole is not detected.
- A normal small circle icon indicates that blackhole is not detected and packet drop anomalies are not present.

Clicking the red alert icon or orange alert icon in the graph filters the corresponding rows in the Alerts table.

Figure 43: Graphs on the Traffic Loss Page



For more information on the Traffic Loss page, see ["Traffic Loss"](#) on page 336.

View Network Topology

IN THIS CHAPTER

- [Network Topology Visualization Overview | 498](#)
- [Network Visualization Options | 501](#)
- [View Network Topology Details | 504](#)
- [Network Table Overview | 515](#)
- [About the Device Tab | 516](#)
- [About the Link Tab | 520](#)
- [About the Site Tab | 524](#)
- [About the SRLGs Page | 526](#)
- [About the Tunnels Tab | 528](#)
- [Add a Tunnel | 535](#)
- [Add Diverse Tunnels | 542](#)
- [About the FlexAlgo Tab | 548](#)
- [Reroute LSPs | 550](#)
- [Add and Remove LSP Delegation | 553](#)
- [Dynamic Topology Workflow | 555](#)
- [Configure a PCE Server | 556](#)
- [Configuring PCEP on a PE Router | 557](#)
- [Configuring Topology Acquisition Using BGP-LS | 559](#)
- [Enable PCEP Security | 560](#)

Network Topology Visualization Overview

Topology visualization enables network administrators to view the complete network topology and its components. By monitoring the network topology, network administrators can monitor network health and plan changes to the topology to keep it functioning optimally.

The topology map in the Topology (**Observability > Topology**) page in Routing Director displays the network topology. The topology map is interactive, and you can customize the map to view the information that you choose. The network information table displayed at the bottom of the page enables you to view detailed information about links, devices, sites, and tunnels in your network and also customize your topology map view.

You can use the Topology page to perform the following tasks:

- **View network topology in the interactive topology map**—The topology map displays the devices, the sites where the devices are located, and the links and tunnels between the devices in your network. The map is a representation of the network, and you can reload the map to refresh the network view. The devices on the map are identified by their hostnames. You can also change the labels to identify devices by their interfaces and IP addresses. You can view the real-time changes in your network if you enable the dynamic topology feature.

For more information, see ["Topology Map" on page 504](#).

- **View live changes in your network**—The topology map displays the live changes in your network. Routing Director establishes a BGP-LS peering session with the devices in your network, and enables you to view live changes in your network even if you do not adopt the devices in your network.

For more information, see ["Topology Menu Bar " on page 509](#).

- **Manage the physical position of devices in the topology map for visual clarity**—If there are many devices clustered together on the map, you can move the devices as per your personal preference or distribute them for easy viewing.

In Device View, the devices and links are displayed as per their coordinates saved in Routing Director. To manage the device positions and map layout, right-click a blank space in the map in Device view, click **Layout**, and select one of the available options. You can also import and export device location information in comma-separated values (CSV) format or GeoJSON format files.

For more information, see ["Topology Map" on page 504](#).

In addition, you can switch between Device or Node View and Cluster View. In cases where devices and links are in proximity in the map, they might overlap with each other and clutter the map. In cluster view, you can collapse devices and links in the topology map into clusters and bundles to view them clearly. Clusters and bundles are aggregated forms of multiple devices and links and reduce the number of network elements visible on the map. Use the vertical topology menu bar on the right to switch between the default device view and cluster view.

For more information, see ["Topology Menu Bar " on page 509](#).

- **View detailed information about the components in your topology in the network information table**—The network information table at the bottom of the page displays detailed information about devices, links, and sites in your topology.

Click the **Device** tab in the network information table to view information about the devices in your topology. View information such as the site where the device is deployed, hostname, IP address, hardware, and severities of the alerts and alarms raised on the device. Alerts and alarms are collectively referred to as events in this topic.

For more information, see ["About the Device Tab" on page 516](#).

Click the **Link** tab in the network information table to view information about the links between your devices. View information such as interface names and the IP addresses of the ingress and egress ports of interfaces connecting the devices. In addition, you can click a link to view information about interface statistics and status of the link.

For more information, see ["About the Link Tab" on page 520](#).

Click the **Site** tab in the network information table to view information about the sites where the devices are located in your topology. View information such as site name, the number of devices in the site, and the severities of the events raised for all devices in the site. In addition, you can also configure new sites from this tab.

For more information, see ["About the Site Tab" on page 524](#).

Click the **Tunnels** tab in the network information table to view information about all label-switched paths (LSPs), also known as tunnels, in your network.

Click the **FlexAlgo** tab in the network information table to view the Flexible Algorithm Definition (FAD) that you have configured on the devices in your network.

For more information, see ["About the FlexAlgo Tab" on page 548](#).

- **View severities of events**—You can view severities of the events in the Device and Site tabs. The Device tab displays the total events raised on each individual device and the Site tab displays the total events generated for all devices in each site. To view more information about the events and to take action to fix the event condition, click the device name to navigate to the **Observability > Troubleshoot Devices > *Device-Name* > Overview** tab.

The Overview tab displays the different network component accordions along with their corresponding health status. If a device has an event, the severity of the event is displayed on the top right of the accordion. Expand the accordion with the event you want to view and perform the necessary action to fix the event condition.

For more information about the *Device-name* page, see ["About the Device-Name Page" on page 461](#).

RELATED DOCUMENTATION

[Network Visualization Options | 501](#)

[Routing Director Overview | 3](#)

Network Visualization Options

IN THIS SECTION

- [Navigation in the Topology Page](#) | 502

To access the Topology page, select **Observability > Network > Topology**.

The Topology page displays the devices, links, sites, and tunnels in your network in graphical and tabular formats. Users with Super User, Network Admin, and Observer roles can monitor the WAN links through the interactive topology map that enables you to customize the view of devices and links across sites. The network table lists the devices, links, sites, and tunnels in your network topology.

The Topology page can be divided into three main components:

- **Interactive Topology Map**—Enables you to view devices, links, Flexible Algorithm (flex algo), Shared Risk Link Groups (SRLGs), and tunnels across sites in your network and customize the map display. Manage the map view, set device positions on the map, customize device and link labels and font size using the right click menu. For more information, see "[View Network Topology Details](#)" on page 504.
- **Topology Menu Bar**—A vertical bar at the top-right corner of the Topology page, which consists of the following:
 - Reset icon—Center the topology map so that it zooms to fit the screen.
 - Plus icon—Zoom in (enlarge) the topology map.
 - Minus icon—Zoom out (reduce) the topology map.
 - Switch to Cluster View/Node View icon—Switch to Cluster View from the default Node View.


In the Node View, all devices are displayed. When devices and links are in proximity in the map, they may overlap with each other and clutter the map. To reduce clutter, you can switch to the Cluster View to collapse devices and links in the topology map into clusters and bundles, respectively. For more information, see "[Cluster View](#)" on page 509.

- **Settings icon**—Configure display properties of the topology map. You can also specify the neighbor (peer) IP address and autonomous system (AS) number so that Routing Director can establish a BGP-LS peering session and enables you to view live changes in your network.
- **Network Table**— Displays detailed information about network devices, links, sites, flex algo, SRLGs, and tunnels. Click the collapsible arrow icon at the bottom-left of the topology map to display or hide the network table.

For more information, see ["Network Table Overview" on page 515](#).

[Figure 44 on page 502](#) shows all the components in the Topology page.

Figure 44: Topology Page



Navigation in the Topology Page

[Table 122 on page 502](#) describes the navigation functions that you can use in the Topology page.

Table 122: Navigation Functions in the Topology Page

Function	Method
Drag and drop	Select an element, drag to the required position on the screen, and then release.
Select an element	Click a link or device to select it.

Table 122: Navigation Functions in the Topology Page *(Continued)*

Function	Method
Select multiple elements	Do one of the following: • Hold down the shift key and left mouse button while dragging the mouse to create a rectangular selection box. All elements within the box are selected. • Hold down the shift key and click multiple items, one at a time.
Zoom in and out 	Do one of the following: • Use the mouse scroll wheel. • Pinch to zoom using the touch pad. • Click the + or - buttons in the topology menu bar.
Reset topology view 	Click the reset icon in the topology menu bar to resize and center the topology map to fit the visible area of the Devices & Links page.
Right-click to access functions	Right-click any blank space in the topology map or on a map element to access context-specific menus.
Collapse/expand pane 	When a left, right, up, or down slider appears at the margin of a pane, you can click it to collapse or expand the pane.
Pin 	Click the pin icon on the top right of an overlay page or a widget in the Topology page to fix it at any place on the screen.
Resize panes	Click and drag any of the pane margins to resize the panes.

View Network Topology Details

SUMMARY

The topology map on the Topology page (**Observability > Network > Topology**) displays the network topology. The map is interactive, which means that you can use the features within the map to customize the map as well as the network information table that is displayed at the bottom of this page.

IN THIS SECTION

- [Topology Map | 504](#)
- [Topology Menu Bar | 509](#)

The map uses a geographic coordinate reference system that enables the following features:

- **Constrained zooming**—The controller checks the coordinates so that the view is constrained to the coordinates of the earth.
- **Repositioning devices according to their geographical coordinates**—By default, each device is positioned in the map according to the geographical coordinates (latitude and longitude) of the site to which the device belongs. If a device is not associated with a site, the device is positioned randomly. You can reposition the devices in the map according to their geographical coordinates if you want to mimic your actual topology in the map.

Topology Map

The topology map displays the network topology. You can right-click a device, link, or blank space in the topology map to access multiple menus. The menus and the options in each of these menus are described in the following table.

Table 123: Options Displayed When You Right Click Blank Space on the Topology Map

Option	Description
Filter in Device Table	You can select a device and filter details of the device in the network information table. Right click a device on the map (device is highlighted with a Yellow circle) and select Filter in Device Table. The Device Table in the Device tab is filtered to display details of only the device that you selected on the map.

Table 123: Options Displayed When You Right Click Blank Space on the Topology Map *(Continued)*

Option	Description
Filter in Link Table	Right click a link on the map and select Filter in Device Table . The Link Table in the Link tab is filtered to display details of only the link that you selected on the map.
Layout Manage the physical position of devices on the topology map using layout options. You can also import and export device information in CSV and GeoJSON formats. The layout options are available in node view.	
Import from	NOTE: Only users with the Super User and Network Admin roles can perform this task. Import information (about the hostname, latitude, longitude, router id, and site information) for all devices in a layout from a comma-separated values (CSV) file or GeoJSON file. For more information about GeoJSON files, see GeoJSON.org. To import the file from your local file system, click Import From and navigate to the folder that contains the CSV or GeoJSON file. Then, click Open to upload the CSV or GeoJSON file. After the file is imported, the devices are automatically repositioned in the topology map according to the coordinates (latitude and longitude) in the imported file. To save the configured coordinates in the server, right-click a blank space in the topology map and select Layout > Set Coordinates from Map.
Export to	NOTE: Only users with the Super User and Network Admin roles can perform this task. Export information (about the hostname, latitude, longitude, router id, and site information) for all devices in a layout as a CSV file or GeoJSON file. Based on the option you select, the CSV or GeoJSON file is automatically downloaded to your local system.

Table 123: Options Displayed When You Right Click Blank Space on the Topology Map *(Continued)*

Option	Description
Set Coordinates from Map	NOTE: Only users with the Super User and Network Admin roles can perform this task. Update the device coordinates in the server according to the current location of devices in the topology map. If you want to reposition the devices according to their coordinates, manually place the devices in the topology map as required. Alternatively, import the coordinates from a CSV or GeoJSON file. Then, use this option to save the current device coordinates in the server.
Reset by Coordinates	NOTE: Only users with the Super User and Network Admin roles can perform this task. Reset the position of devices in the topology map according to the device coordinates, as retrieved from the server. If a device does not have configured coordinates, the device is automatically repositioned according to the coordinates of the site to which the device belongs.
Toggle Background Map	By default the topology map view loads the world map. Right click and select Toggle Background Map to turn off the world map background. If you repeat the action, Routing Director loads the map again.
Save to Default Map Layout	Saves the current map layout as the default map layout. For example, if you turn off the world map background and select Save to Default Map Layout, Routing Director saves this settings. Next time when you revisit the Topology page, the topology map will not load the world map.

Table 123: Options Displayed When You Right Click Blank Space on the Topology Map *(Continued)*

Option	Description
Label Size	• Select a Label Size: Select one of the following values as the font size for the device and link labels: • 10 • 12 • 14 • 16 • 18 • 20
Device Label	Select one of the following options to label the devices in the topology map: • Hostname • Hostname, SID • IP Address • OS Version • SID • Hide Label—Hides all the labels for the devices in the topology map • IP, SID

Table 123: Options Displayed When You Right Click Blank Space on the Topology Map *(Continued)*

Option	Description
Link Label	Select one of the following options to label the links in the topology map: • Hostname A::Z • Interface A::Z • IP A::Z • IP, SID A::Z • TE Metric A::Z • Interface Util A::Z • ISIS1 Metric A::Z • ISIS2 Metric A::Z • OSPF Metric A::Z • RSVP Bandwidth A::Z • RSVP Util A::Z • RSVP Live Util A::Z • Measured Delay A::Z • SID A::Z • TE Admin Group A::Z • Hide Label • Hide Label—Hides all the labels for the links in the topology map
Reload Network	Reloads the network, and updates the displayed topology map.

Topology Menu Bar

IN THIS SECTION

- [Cluster View | 509](#)
- [Topology Settings Options | 510](#)

The topology menu bar is the vertical bar at the top-right corner of the Topology page, which consists of the following:

- Reset icon—Center the topology map so that it zooms to fit the screen.
- Plus icon—Zoom in (enlarge) the topology map.
- Minus icon—Zoom out (reduce) the topology map.
- Switch to Cluster View/Node View icon—Switch to Cluster View from the default Node View.

In Node View, all devices and links are displayed as is. In cases where devices and links are in proximity in the map, they might overlap with each other and clutter the map. In Cluster View, you can collapse devices and links in the topology map into clusters and bundles, respectively, to reduce clutter. For more information, see "[Cluster View](#)" on page 509.

- Settings icon—Configure display properties of the topology map, and specify the neighbor (peer) IP address and autonomous system (AS) number for viewing the live changes in your network. For more information, see "[Topology Settings Options](#)" on page 510.

Cluster View

By default, the GUI displays the topology map in the Node view. In this view, all devices and links are displayed as is. In cases where devices and links are in proximity in the map, they might overlap with each other and clutter the map. To reduce clutter, **select the Switch to Cluster View** option in the topology menu bar. The overlapping devices and links automatically collapse into clusters and bundles, respectively. Isolated devices and links remain as is. The clusters and bundles reduce visual clutter in the topology map and aggregate data, enabling you to view the network better, especially in case of large-scale networks with many devices and links.

To return to the default view, **select the Switch to Node View** option in the topology menu bar.

Cluster View has the following features:

- Each cluster is represented by a circle. The number in each circle indicates the number of devices in the cluster. Similarly, each bundle is represented by a thick line. The number on the line indicates the number of links in the bundle.
- When you double-click a cluster, the topology map zooms in to expand the cluster into its child devices. When you double-click a bundle, the bundle expands to display individual links. To collapse the links back into a bundle, double-click the underlay hull.

A hull is a visual representation that is drawn behind a bundle to allow the user to expand or collapse curved lines (that is, the links in the topology map).

- You can drag clusters to reposition them in the topology map. As a result, the devices in the cluster and the links connecting these devices are also repositioned.
- If two or more devices in a cluster are directly connected through a common link, the circle representing the cluster displays a colored outline. The color of the outline is the same as the color configured to denote the highest utilization for the link interconnecting those devices.

To identify interconnected devices in a cluster, double-click the cluster and zoom in to the next level.

Topology Settings Options

[Table 124 on page 511](#) describes the tabs on the Topology Settings pane.

Table 124: Topology Settings Options

Tab	Description
Nodes	You can perform the following tasks: • Click the Label drop down to choose the label that you want to be displayed on the topology map. The following options are available: • None—No labels are displayed on the topology map. • Hostname—The hostname of the node is displayed on the topology map. • IP Address—The IP address of the node is displayed on the topology map. • OS Version—The OS version of the node is displayed on the topology map. • Hostname, SID—The hostname and SID of the node is displayed on the topology map. • IP, SID—The IP address of the node and the SID is displayed on the topology map. • SID—The SID of the node is displayed on the topology map. • After you select a label, you can enable the following toggle buttons: • Background Shadow—To view the background shadow of node labels. • Hide Pseudo Node Label—To hide the pseudo node label • Show only Favorite Labels—To view only favorite node labels • Hide Isolated Nodes—To hide isolated nodes • Hide Pseudo Nodes—To hide pseudo nodes • Show Site in a Circle—To display site in a circle.

Table 124: Topology Settings Options (*Continued*)

Tab	Description
Links	You can perform the following tasks: - Click the Label drop down to choose the label that you want to be displayed on the topology map. The following options are available: - None - Hostname A::Z - Interface A::Z - IP A::Z - IP , SID A::Z - TE Metric A::Z - ISIS1 Metric A::Z - ISIS2 Metric A::Z - OSPF Metric A::Z - Interface Util A::Z - RSVP Bandwidth A::Z - RSVP Util A::Z - RSVP Live Util A::Z - Measured Delay A::Z - SID A::Z - TE Admin Group A::Z - You can select the network element that you want to be displayed in the color legend (bottom right corner of the Topology map). You can select one of the following network elements: - RSVP Utilization

Table 124: Topology Settings Options *(Continued)*

Tab	Description
	• RSVP Live Utilization • Interface Utilization • Measured Interface Delay • Click the toggle button to: • Show or hide the link down markers • Draw down links as a dashed line • Draw a link line width in proportion to the interface speed • Draw parallel links as a curve: Toggle the button to draw parallel links between two nodes as a curve, so that the parallel links do not overlap and appear separately (as curves) on the topology map. If the parallel links between two nodes are drawn as straight lines, they would overlap on the topology map (as a bundle). • Wrap links as great arcs: Distinguishes links that would have to wrap around the world map. • Hide Partially Visible links: Hides any link whose end nodes are outside the visible area. This is useful for focusing on a subset of a large network.

Table 124: Topology Settings Options (*Continued*)

Tab	Description
General	You can perform the following tasks: - Click the Enable Animation while calculating Topology Layout toggle button to enable animation while calculating the topology layout. - Select a Label Size: Select one of the following values as the font size for the node and link labels: 10 12 14 16 18 20 - Apply Opacity effects. Move the Opacity Effects slider to select the percent opacity for topology map elements that are not highlighted. - Select the color for nodes and links.
Dynamic Topology	Routing Director establishes a BGP-LS peering session with the devices in your network, and enables you to view live changes in the network even if the devices are not managed by Routing Director. To view live updates in topology, you need to do the following: - Specify the IP address of the BGP-LS peer and autonomous system (AS) number on the Topology Settings page (Observability > Network > Topology > Topology Menu Bar > Settings icon). - Manually refresh the Topology page or you need to right-click on the topology map and select Reload Network. In addition, on the topology map, you can view the latest operational status of devices, links, sites, and tunnels.

Network Table Overview

The network table at the bottom of the Topology page displays detailed information about devices, links, sites, and tunnels in their respective tabs. To access the network table, click **Observability > Network > Topology**. You can perform the following common functions from the device tab, the link tab, the sites tab, and the tunnels tab on the Topology page.

- **Hide/Show the network table**—To display or hide the network information table, click the collapsible arrow icon present in the top left-corner of the table.
- **Download**—To download the data displayed in the selected tab to your local system, click **Download**. The data is downloaded to your local system as a comma-separated values (CSV) file.
- **Filter Entries Using Criteria**—To filter the table entries by adding new filtering criteria, hover over the **Filter** (funnel) icon and select **Add Filter**. On the **Add Criteria** page that appears, select the filtering criteria from the **Field** and **Condition** list, and enter the text to be compared in the **Value** field. Then, click **Add**.

The filtered table entries are listed and the filter criteria name is displayed above the table column names.

To remove the filtering criteria, click the (X) icon (next to the filter name).



NOTE:

- You can add multiple filtering criteria. Once you add the multiple filtering criteria, select the **And** condition to display the entries matching all the filtering criteria or the **Or** condition to display the entries matching any one of the filtering criteria.
- **Quick filter**: Once you have added all the filtering criteria, click **Save** to save a particular criterion or multiple criteria for future use.

On the Save Filter page that appears, enter a name for the filter. Optionally, toggle the **Set as Default** button if you want to use this filtering criteria by default, and click **OK**.

The saved filters are displayed under **Quick Filters** when you hover over the Filter (funnel) icon. You can then apply these saved filters to the table entries.

- **Show/Hide Columns**—Choose to show or hide one or more columns in the table on each tab.

Hover over the vertical ellipsis icon and select **Show/Hide Columns**. In the list that appears, select the *Column-Name* check boxes corresponding to the columns you want to display in the table.

Only the selected columns are displayed in the table.

- **Reset Preference**—Resets the displayed columns to the default set of columns in the table for each tab and reloads the topology map.

Hover over the vertical ellipsis icon and select **Reset Preference**.

Only the default columns are displayed in the table. Routing Director also reloads the topology map to the default view if you reposition devices without saving the coordinates or filter links.

The network table has the following tabs:

- **Device**—View details about the devices in the network. For more information, see ["About the Device Tab" on page 516](#).
- **Link**—View details about the links in the network. For more information, see ["About the Link Tab" on page 520](#).
- **Sites**—View details about the sites where you deployed devices in the network. For more information, see ["About the Site Tab" on page 524](#).
- **Tunnels**—View details about the tunnels in the network. For more information, see ["About the Tunnels Tab" on page 528](#).
- **FlexAlgo**—View details about the Flexible Algorithm (flex algo) IDs in your network. For more information, see ["About the FlexAlgo Tab" on page 548](#).

About the Device Tab

IN THIS SECTION

- [Tasks You Can Perform | 517](#)

You can view information (such as the hostname, IP address, and type) about the devices in the network in the Device tab of the network table (**Observability > Network > Topology > Device**). [Table 125 on page 518](#) provides information on the fields in the Device tab.

You can also perform various actions on the devices from this tab. To perform common actions such as filtering using advanced filter criteria, downloading device details, and resetting preferences, see ["Network Table Overview" on page 515](#).

Tasks You Can Perform

- From the **More** list, you can perform the following tasks:



NOTE: You can also right-click a device row in the table to view the same options.

- View details about a device—To view details (such as the properties configured for a device and its interfaces):
 - Hover over a device row in the table and click the **Details** icon. On the **Device Device-Name** page that appears, you can view the Details tab and the Interfaces tab.
 - Select the device on the table and click **More > Show Detail**. On the **Device Device-Name** page that appears, you can view the Details tab and the Interfaces tab.

On the Details tab, you can view general device properties such as, model, device type, MAC address, and the serial number. You can also view site properties such as site name, address, country, and time zone where the device is deployed. On the Interfaces tab, you can view all the interfaces of the device, the administrative and operational statuses, the speed of the physical interfaces (in Mbps), and the duplex mode. The **Device Device-Name** page is moveable, so you can pin it anywhere on the screen.



NOTE: Routing Director populates the speed column for physical interfaces and for interfaces with Operational Status Up.

- Filter a device in the topology map—To display only the selected device in the topology map, select **Filter Selected device**. The Devices & Links page reloads to display all devices and links on the topology map. To undo the filter, right-click a blank space on the topology map and click **Reload Network**.
- Zoom in on a device—To zoom in on a device in the topology map, **select** the device and click **Zoom to Selected device**. The topology map is enlarged to zoom in on the selected device.
- Assign a device to a site—You can select devices and add them to a site. See ["About the Troubleshoot Devices Page" on page 454](#) for steps to assign devices to sites.
- Back up device configuration—You can back up device configuration. In the event of a device failure, you can use the backed up configuration to restore the device configuration. See ["About the Troubleshoot Devices Page" on page 454](#) for the steps to back up a device configuration.
- Reboot—You can reboot the device after making configuration changes or commit changes in a planned maintenance event. See ["About the Troubleshoot Devices Page" on page 454](#) for the steps to reboot the device.

- Upgrade—You can upgrade the current software image on the device to the latest version. See ["About the Troubleshoot Devices Page" on page 454](#) for the steps to upgrade the software image.

Table 125: Fields on the Device Table

Fields	Description
Hostname	Lists the physical and virtual hosts (devices) in your network.
Severity	Displays the highest severity of the events on the device. When you hover over the severity, a pop-up displays the total number of events for all severity levels. The severity levels are: • Urgent Action Needed or critical • Action Needed or major • Being Monitored or minor If there are no events, you can see a green ✓ icon, indicating a healthy device status.
IP Address	Displays the management IPv4 address of the physical and virtual devices in your network.
Site	Displays the geographical site to which you added the device.
Model	Displays the model of the device. For example, ACX-7100-32C.
Serial Number	Displays the serial number of the device.
Latitude	Displays the latitude value of a device. Latitudes range from -90 to 90. Positive values of latitude are north of the equator and negative values (preceded with a minus sign) are south of the equator.

Table 125: Fields on the Device Table *(Continued)*

Fields	Description
Longitude	Displays the longitude value of a device. Longitudes range from -180 to 180. Positive longitudes are east of the Prime Meridian and negative values (preceded with a minus sign) are west of the Prime Meridian.
OS Version	Version of the OS that is currently installed on the device.
PCEP Status	Display whether path computation element protocol (PCEP) is configured on a device or not.
PCEP IP	Displays the IP address of the PCEP.
SID	Displays the segment identifier (SID) if you have configured segment routing.
SR	Display whether segment routing (SR) is configured on a device or not.
SRLG	Displays information about shared risk link group (SRLG).
FAD Definition	Displays the flexible algorithms definition (FAD) ID. Hover over the FAD ID link to view the flex algo configuration.
FAD Participation	Displays the list of FAD IDs that the device is part of.

RELATED DOCUMENTATION

Network Visualization Options | 501

About the Link Tab

IN THIS SECTION

- [Tasks You Can Perform | 520](#)

You can view information about the links in the network, in the Link tab of the network table (**Observability > Network > Topology > Link**). You can also perform various actions on the links from this tab. To perform common actions such as filtering using advanced filter criteria, downloading link details, and resetting preferences, see ["Network Table Overview" on page 515](#).

Tasks You Can Perform

- From the **More** list, you can perform the following tasks:



NOTE: You can also right-click a link to view the same options.

- View details about a link—To view details such as the site name, country, site id, organization, id, and time zone:
 - Hover over a link row in the table and click the **Details** icon. The **Link - Device A to Device Z** appears. On the **Link - Device A to Device Z** page that appears, you can view the Interface Stats tab and the Details tab.
 - Select the link on the table and click **More > Show Detail**. On the **Link - Device A to Device Z** page that appears, you can view the Interface Stats tab and the Details tab.

On the Interface Stats tab, you can view the device name, interface name, and management IP address of device A and device Z. On the Details tab, you can view the node id associated with the interfaces, interface IPv4 address at the two nodes, interface name, interface bandwidth, live status, and the operational status of the interface. The **Link - Device A to Device Z** page is movable so, you can pin it anywhere on the screen.

- Filter links on the topology map—To display only the selected link in the topology map, select **Filter Selected Link**. To undo the filter, right click a blank space on the topology map and click **Reload Network**.
- Zoom in on a link—To zoom in on a link in the topology map, select the link and click **Zoom to Selected Link**. The topology map is enlarged to zoom in on the selected link.

- Configure table settings—Select **Table Settings** to specify the link utilization format (**Decimal** or **Percentage**) for the link utilization parameters in the specified format.
- Download link-related information—Click **Download** to download detailed information about all the existing links in the topology in CSV format.
- View historical data—Routing Director enables you can view the historical data. Historical data provides you valuable insights into past trends and patterns.



NOTE: For Routing Director to collect historical data, you need to do the following:

1. You need to enable the collection of statistics.
 - a. To collect historical data for interface traffic, you need to enable traffic engineering while creating a device profile. For more information, see ["Add a Device Profile" on page 208](#).
 - b. To collect historical data for interface delay, you need to enable link delay measurements. For more information, see <https://www.juniper.net/documentation/us/en/software/junos/is-is/topics/topic-map/enable-link-delay-advertise-in-is-is.html>
2. Enable dynamic topology so that you can view real-time changes in your network. For more information, see ["Dynamic Topology Workflow" on page 555](#).

You can view the historical data for the following:

- Interface Traffic—View a graphical representation of the link traffic (in bps) for interface A and Z based on the time range selected from the **Period** list. A **Link Traffic** pop-up appears, which you can pin anywhere on the screen using the pin icon. From the **Period** list, you can select an option to view the traffic data for Previous 3 Hours, 1 day, 1 Week, or Custom time ranges.

You can choose to auto refresh the data or manually refresh the data.

- Interface delay—View a graphical representation of the link delay (in milliseconds) for interface A and Z based on the time range selected from the **Period** list. A **Link Delay** pop-up appears, which you can pin anywhere on the screen using the pin icon. From the **Period** list, you can select an option to view the delay data for Previous 3 Hours, 1 day, 1 Week, or Custom time ranges.

You can choose to auto refresh the data or manually refresh the data.

Table 126: Fields on the Link Table

Fields	Description
Device A	Displays the device at which traffic enters.
Device Z	Displays the device at which traffic exits.
Status	Displays the status of the link.
IP A	Displays the IPv4 address of the interface from which device A sends traffic. Routing Director displays the IPv4 address of IP A based on device A's active configuration. If the IPv4 addresses of IP A and IP Z are in the same subnet, Routing Director forms an interface between devices A and Z.
IP Z	Displays the IPv4 address of device Z interface that receives traffic. Routing Director displays the IPv4 address of IP Z based on device Z's active configuration. If the IPv4 addresses of IP A and IP Z are in the same subnet, Routing Director forms an interface between devices A and Z.
Interface A	Displays the interface name of device A.
Interface Z	Displays the interface name of device Z.
Bandwidth A	Displays the bandwidth (in bps) of the ingress device.
Bandwidth Z	Displays the bandwidth (in bps) of the egress device.
TE Metric A	Displays the traffic engineering-related metrics for the device A.
TE Metric Z	Displays the traffic engineering-related metrics for the device B.
Admin Group AZ	Displays the admin group selected for the traffic to traverse between device A and device Z.

Table 126: Fields on the Link Table *(Continued)*

Fields	Description
Admin Group ZA	Displays the admin group selected for the traffic to traverse between device Z and device A.
Interface Util AZ	Displays the interface utilization between device A and device Z.
Interface Util ZA	Displays the interface utilization between device Z and device A.
RSVP BW AZ	Displays the RSVP bandwidth between device A and device Z.
RSVP BW ZA	Displays the RSVP bandwidth between device Z and device A.
RSVP Util AZ	Displays the RSVP utilization between device A and device Z.
RSVP Util ZA	Displays the RSVP utilization between device Z and device A.
Measured Delay AZ	Displays the delay metrics between device A and device Z.
Measured Delay ZA	Displays the delay metrics between device Z and device A.
SID A	Displays the SID label for the ingress device.
SID Z	Displays the SID label for the egress device.
ISIS1 AZ	Displays the ISIS (area 1) metric between device A and device Z.
ISIS1 ZA	Displays the ISIS (area 1) metric between device Z and device A.

Table 126: Fields on the Link Table *(Continued)*

Fields	Description
ISIS2 AZ	Displays the ISIS (area 2) metric between device A and device Z.
ISIS2 ZA	Displays the ISIS (area 2) metric between device Z and device A.
OSPF Metric AZ	Displays the OSPF metric between device A and device Z.
OSPF Metric ZA	Displays the OSPF metric between device Z and device A.
SRLG	Displays the SRLGs (or SRLG IDs) associated with the link. This information helps identify the links that share the same SRLGs.

RELATED DOCUMENTATION

| [Network Visualization Options](#) | 501

About the Site Tab

IN THIS SECTION

● [Tasks You Can Perform](#) | 525

You can view information about the sites in the network, in the Site tab of the network table (**Observability > Network > Topology > Site**). Sites are the physical locations that host the devices within an organization's network. You can also perform various actions on the sites from this tab.

Tasks You Can Perform

- From the **More** list, you can perform the following tasks:
 - View details about a site—To view details such as site name, site id, country, and time zone:
 - Hover over a site row in the table and click the **Details** icon.
 - Select the site on the table and click **Show Detail** from the **More** drop-down list.

On the **Site *Site-Name*** page that appears, you can view the number of events by severity levels, address, country, site id, name, organization id, and time zone of the site. The page is movable so, you can pin it anywhere on the screen.

- Add a Site—To add a site, click **Add Site** from the **More** drop-down list. The Create Site page appears. See ["Add Sites" on page 90](#) for the procedure to create a site.
- To perform common actions such as filtering using advanced filter criteria and resetting preferences, see ["Network Table Overview" on page 515](#).

Table 127: Fields on the Site Table

Fields	Description
Id	Displays the unique id for a site.
Name	Displays the name of the site.
Severity	Displays the highest severity of the events on the device. When you hover over the severity, a pop-up displays the total number of events for all severity levels. The severity levels are: • Urgent Action Needed or critical • Action Needed or major • Being Monitored or minor If there are no events, you can see a green ✓ icon, indicating a healthy device status.
Device Count	Displays the number of devices in each site.
Country	Displays the country where the site is located.

Table 127: Fields on the Site Table *(Continued)*

Fields	Description
Timezone	Displays the time zone of the site.
Address	Displays the complete address of sites.

RELATED DOCUMENTATION

[About the Device Tab | 516](#)

[About the Link Tab | 520](#)

About the SRLGs Page

IN THIS SECTION

- [Tasks You Can Perform | 527](#)
- [Field Description | 527](#)

To access the SRLGs tab, click **Observability > Network > Topology**.

A Shared Risk Link Group (SRLG) is a set of links sharing a common resource. These links share the same risk of failure and are therefore grouped to belong to the same SRLG.

On this tab, you can view the list of SRLGs in your network. You can also view the number of links associated with an SRLG and determine the links that share the same SRLG, which is crucial for network redundancy and planning.

You need to establish BGP-LS session between Routing Director and the router so that the router can advertise SRLG-related information configured in your network. Otherwise, you cannot view the list of SRLGs configured in your network on the SRLGs tab.

Tasks You Can Perform

- View details about the SRLG by clicking the details icon when you hover anywhere over the SRLG row or click **More > Show Detail**.
- Filter the data displayed in the table—Click the filter icon (funnel) and select whether you want to show or hide advanced filters. You can then add or remove filter criteria, save criteria as a filter, apply or clear filters, and so on. The filtered results are displayed on the same page.
- Show or hide columns in the table or reset page preferences, using the vertical ellipsis menu.
- Sort, resize, or re-arrange columns in a table (grid).

Field Description

[Table 128 on page 527](#) lists the fields on the SRLGs tab.

Table 128: Fields on the SRLGs Tab

Fields	Description
ID	Displays the SRLG ID.
Source	Displays the source (BGP-LS) from which the SRLG-related information is fetched.
Links	Displays the number of link associated with the SRLG ID
Link Names	Displays the list of link names that are associated with the SRLG ID.

RELATED DOCUMENTATION

| [Network Visualization Options](#) | 501

About the Tunnels Tab

IN THIS SECTION

- [Tasks You Can Perform](#) | 528

On the Tunnels tab of the network table (**Observability > Network > Topology > Tunnels**), you can view the information about label-switched paths (LSPs) in your network.

Routing Director uses Path Computation Element Protocol (PCEP) sessions to discover the tunnels that run in your network.



NOTE: For Routing Director to discover tunnels in your network:

1. A PCEP session must be established between Routing Director and the devices.

To establish a PCEP session, you must specify a VIP address either at the time of installing Routing Director or post-installation of the cluster. For more information, see ["Configure a PCE Server" on page 556](#).

To establish PCEP on a PE router, see ["Configuring PCEP on a PE Router" on page 557](#).

To configure BGP-LS topology acquisition on Routing Director, see ["Configuring Topology Acquisition Using BGP-LS" on page 559](#).

2. You must have adopted or onboarded the device to Routing Director; otherwise, the tunnel-related information is not displayed on the Tunnels tab.

Tasks You Can Perform

- Hide unrelated nodes of a tunnel—To hide unrelated nodes of a tunnel, select one or more tunnels and enable the **Hide unrelated** toggle button. The topology map limits the display to only the nodes that are related to the selected tunnel.
- Provision tunnels—You can provision the following:
 - Provision a tunnel—Click **Provisioning** and select **Tunnel** to provision a tunnel. For more information, see ["Add a Tunnel" on page 535](#).

- Provision secondary or standby tunnels—If you want to provide an alternate route in case the primary route fails, you can add additional secondary or standby tunnels for a tunnel. Select the tunnel and click **Provision Secondary/Standby**.

The Add Tunnel page appears, where you can select a secondary or standby tunnel for the existing tunnel. After the secondary or standby tunnel that you created is provisioned, you can see it in the network information table and in the topology map.



NOTE: You can provision a secondary or standby tunnel only for RSVP tunnels. The **Provision Secondary/Standby** option is available only for PCC-initiated and PCC-controlled tunnels.

Routing Director supports the discovery, control, and creation of primary and protection LSPs. Primary LSPs provide the primary (preferred) route for traffic flows, while protection LSPs provide an alternate route if the primary route fails.

Protection LSPs are of two types: standby LSPs and secondary LSPs. The tunnel ID, source node, destination node, and IP address of a secondary or standby LSP are identical to that of the primary LSP. However, secondary and standby LSPs have the following differences:

- A secondary LSP is not signaled until the primary LSP fails.
- A standby LSP is signaled regardless of the status of the primary LSP.

When you configure a protection LSP from the GUI, you must have a primary LSP of the same control type (PCC-controlled, PCC-delegated, or PCE-initiated) available before you configure a protection LSP.



NOTE: By default, a protection LSP uses the bandwidth, setup priority, and hold priority values of the primary LSP. However, each protection LSP can be configured to use values different from the primary LSP.

- Provision Diverse Tunnels—Click **Provisioning** and select **Diverse Tunnels** to implement tunnel diversity in your network. For more information, see ["Add Diverse Tunnels" on page 542](#).
- Trigger tunnel optimization—Routing Director ensures that LSPs are automatically routed through the most optimal paths. However, if you want Routing Director to recalculate an optimal path, click **Provisioning** and select **Trigger Tunnel Optimization**. For information on how Routing Director automatically reroutes LSPs see ["Reroute LSPs" on page 550](#).
- Add or Remove delegation—To delegate one or more PCC-controlled LSPs to the PCE or return control of delegated LSPs to the PCC, select **Configure Delegation**.

For more information, see ["Add and Remove LSP Delegation" on page 553](#).

- View historical data—Routing Director enables you to view historical data. Historical data provides you valuable insights to past trends and patterns.

For Routing Director to collect historical data, you need to do the following:

1. Enable the collection of statistics.
 - To collect historical data for tunnel traffic, you must enable traffic engineering while creating a device profile. For more information, see ["Add a Device Profile" on page 208](#).
 - To collect historical data for tunnel delay, you must enable link delay measurements. For more information, see [How to Enable Link Delay Measurement and Advertising in IS-IS](#).
2. Enable dynamic topology so that you can view real-time changes in your network. For more information, see ["Dynamic Topology Workflow" on page 555](#).

You can view the historical data for the following:

- Tunnel traffic—To view traffic and bandwidth for a tunnel in graphical form, select the tunnel and click **Tunnel Traffic**. In the Tunnel Traffic page that appears, select the period for which you want to view the data. You can view data for the previous 3 hours, the previous day, the previous week, or choose from a custom time range by specifying the start and end dates and the time.



NOTE: You can view the tunnel traffic if the device has Junos OS Release 22.X or later releases.

If you are using ACX Series device, then you can view tunnel traffic only if you have installed Junos OS Evolved Release 23.4R2-S2.1.

- Tunnel delay—Tunnel delay is the sum of all link delays in the tunnel path. At any given time, the Path Computation Element (PCE) is aware of the paths of all tunnels in the network. Periodically, the PCE uses the reported link delays to compute the end-to-end tunnel delay.

To view tunnel delay in graphical form, select the tunnel and click **Delay**. In the Tunnel Delay page that appears, select the period for which you want to view the data. You can view data for the previous 3 hours, the previous day, the previous week, or choose from a custom time range by specifying the start and end dates and the time.

You can choose to refresh the data automatically or manually.



NOTE: We do not support traffic statistics collection for Segment Routing (SR) LSPs.

- View events for tunnels—To view historical events (such as actions performed on the tunnel and bandwidth changes) for a tunnel for a specific time range, select the tunnel and click **View > Event History**.

On the Events page, you can:

- View a graphical representation of bandwidth changes for an LSP to analyze trends over time. Select the period for which you want to view the bandwidth-related data. You can view data for the previous 2 hours, 4 hours, 8 hours, 16 hours, 24 hours, 1 week, or choose from a custom time range by specifying the start and end dates and the time.
- View changes to label-switched path (LSP) attributes and events in the table. The Timestamp column helps you to identify the time at which a change has been made to a tunnel.
- View changes to an LSP on the Topology map. You can use the **More > Show Path Changes** option to view highlights on the Topology map for any changes to the LSP path, making it easier to track and analyze routing changes.
- Download comma-separated values (CSV) file that includes information on LSP event history to your local system by clicking the **More > Download** option.



NOTE: You can track event history for all types of LSPs only if the router reports the LSP to Router Director using PCEP protocol. If PCEP is not enabled on the router, then the LSPs are not reported and the LSP event history is not displayed. The default retention policy for LSP event history data is three days.

- Download—To download tunnels-related data to your local system, click **Download**. The data is downloaded to your local system as a comma-separated values (CSV) file.
- From the More list, you can perform the following tasks:
 - View details of a tunnel—To view the details of a tunnel, you can do one of the following:
 - Select a tunnel on the table and click **More > Show Details**. The **Tunnel - Tunnel-Name** page appears. You can view the To and From traffic-related information.
 - Hover over a tunnel row in the table and click the **Details** icon. The **Tunnel - Tunnel-Name** page appears. You can view the To and From traffic-related information.
 - Reload—To view the latest changes in the network map, click **Reload**.
- To perform common actions such as filtering using advanced filter criteria and resetting preferences, see "[Network Table Overview](#)" on page 515.

Table 129: Fields on the Tunnel Table

Fields	Description
Name	Displays the name of the tunnel.
Device A	Displays the device at which traffic enters.
Device Z	Displays the device at which traffic exits.
IP A	Displays the IPv4 address of the interface from which device A sends traffic. Routing Director displays the IPv4 address of IP A based on device A's active configuration. If the IPv4 addresses of IP A and IP Z are in the same subnet, Routing Director forms an interface between devices A and Z.
IP Z	Displays the IPv4 address of device Z interface that receives traffic. Routing Director displays the IPv4 address of IP Z based on device Z's active configuration. If the IPv4 addresses of IP A and IP Z are in the same subnet, Routing Director forms an interface between devices A and Z.
Vendor	Displays the vendor name of the device.
Bandwidth	Displays the traffic rate (bps) associated with the LSPs, fast reroutes, or multiclass LSPs.
Color	Displays the color associated with the segment routing (SR) policy. A segment routing (SR) policy uses the color community value to steer traffic flows on to a tunnel.

Table 129: Fields on the Tunnel Table *(Continued)*

Fields	Description
Admin Group Include All	Displays a list of all admin groups that are created for LSPs. Admin groups are created to implement a variety of policy-based LSP setups.
Admin Group Include Any	Displays a list of the admin groups that are created for LSPs. Admin groups are created to implement a variety of policy-based LSP setups.
Metric	Displays the metric value of an LSP. The LSP metric is used to indicate the ease or difficulty of sending traffic over a particular LSP. Lower LSP metric values (lower cost) increase the likelihood of an LSP being used. Conversely, high LSP metric values (higher cost) decrease the likelihood of an LSP being used.
Path Type	Displays the type of LSP path—primary, secondary, or standby.
Path Name	Displays the user-defined name of the LSP path.
Type	Displays the type of the LSP—RSVP or SR.
Op Status	Displays the operational status of the LSP—Failed, pending, unknown, live.
Signaling address	Displays the signaling address Egress Peer Engineering (EPE).
Path Intent	Displays the name of the path intent, if any.

Table 129: Fields on the Tunnel Table *(Continued)*

Fields	Description
Control Type	Displays the control type.
Setup	Displays the setup priority of the LSP. The setup priority ranges from 0 (highest) through 7 (lowest).
Hold	Displays the RSVP hold priority for the tunnel traffic. The hold priority ranges from 0 (highest) through 7 (lowest).
Record Route	Displays whether Routing Director should actively record the path by sending the record route object (RRO). The RRO is used to record the path that the LSP traverses. It includes the IP address and router/node ID of each RSVP node in the path.
Last Update	Displays the last updated time of the LSP.
Controller Status	Displays the controller status—No path found, Reschedule in <i>X</i> minutes.
RRO Hop Count	Displays the hop count of the record route object (RRO). The RRO keeps track of the actual path the LSP is traversing.
Comment	Displays the comment that the user has added at the time of tunnel creation.
Admin Group Exclude	Displays the admin group that is excluded. The tunnel traverses links that do not include any of the admin groups listed here.
Planned Metrics	Displays the static tunnel metric. The PCE uses this metric to route the tunnel instead of allowing the router to choose a path

Table 129: Fields on the Tunnel Table *(Continued)*

Fields	Description
FlexAlgo ID	Displays the flex algo ID that is associated with the LSP.
Diversity Group	Displays the name of the diversity group to which this tunnel belongs
Diversity Level	Displays the diversity level (Link, SRLG, or site) of the tunnel.

Add a Tunnel

Routing Director allows you to add label-switched paths (LSPs) by using Path Computation Element Protocol (PCEP). You can create RSVP and segment-routed (SR) LSPs, and view the LSPs on the topology map.

LSPs are referred to as tunnels in the Routing Director GUI.

To provision a tunnel:

1. Navigate to **Observability > Network > Topology**.

The Topology page is displayed with the topology map at the center and the network information table at the bottom of the page.

2. Navigate to the Tunnels tab.
3. Click the **Provisioning** drop down and select **Tunnel**.

The Add Tunnel Page appears.

4. Complete the configuration on each tab according to the guidelines in [Table 130 on page 536](#).



NOTE: Fields marked with an asterisk (*) are mandatory.

5. (Optional) From any tab, click **Preview Path** at the bottom of the page to view the path on the topology map.
6. Click **Add** to add the tunnel.

A confirmation message appears on the top of the page, indicating that a provision tunnel request was successfully created.

7. Verify that the LSPs are configured correctly. In the Tunnels tab of the Network Information table:
 - a. Verify that the Operation Status field displays Active.
 - b. Select the LSP and check the path of the LSP on the topology map.
 - c. Check the tunnel traffic and tunnel delay by selecting the LSP and clicking **View > Tunnel Traffic** and **View > Delay** respectively.

Table 130: Fields on the Add Tunnel Page

Field	Description

PROPERTIES

Provisioning Method	From the list, select one of the following methods to be used to provision the tunnel: • PCEP (Path Computation Element Protocol)—The path computation element (PCE) initiates the tunnel and the associated configuration statements do not appear in the router configuration file. Upon provisioning, this tunnel is added as a PCE-initiated tunnel. This is the default provisioning method. • NETCONF—The tunnel is statically provisioned and the associated configuration statements appear in the router configuration file. Upon provisioning, this tunnel is added as a device-controlled tunnel. NOTE: • • For Cisco IOS-XR routers, NETCONF-based tunnel provisioning has the same capabilities as PCEP-based tunnel provisioning. • When provisioning tunnels by using NETCONF one at a time, the provisioning order might be sent before the response to a previous provisioning order is received. The second order might not have the correct bandwidth allocation information and the PCE might not be able to provide ECMP. We recommend provisioning multiple tunnels through NETCONF in one operation (bulk provisioning) in order to avoid this issue.

Table 130: Fields on the Add Tunnel Page (*Continued*)

Field	Description
Provision Type	From the list, select the type of tunnel that you want to provision: • RSVP • SR (segment routing) • SRv6
Name	Specify a unique name for the tunnel. You can use any number of alphanumeric characters, period, hyphens, colons, and underscores. NOTE: If you are adding multiple parallel tunnels that will share the same design parameters, the name you specify here is used as the base for automatically naming those tunnels. See the Count and Delimiter fields in the Advanced tab for more information.
Device A	From the list, select the device that you want to use as the ingress node.
Device Z	From the list, select the device that you want to use as the egress node.
Path Type	From the list, select primary, secondary, or standby as the path type.
Path Name	Specify the name for the path. This field is available only for primary tunnels with RSVP provisioning type, and for all secondary and standby tunnels.

Table 130: Fields on the Add Tunnel Page (*Continued*)

Field	Description
Planned Bandwidth	Specify the planned bandwidth (along with valid units, with no space between the bandwidth and units) for the tunnel. If you specify a value without units, bps is automatically applied. Valid units are: • B or b • M or m • K or k • G or g Examples: 50M, 1000b, 25g.
Planned Metric	Specify the static tunnel metric. The PCE uses this metric to route the tunnel instead of allowing the router to choose a path.
Setup	Specify the setup priority for the tunnel traffic. Priority levels range from 0 (highest priority) through 7 (lowest priority). The default is 7, which is the standard MPLS tunnel definition in Junos OS. Based on the setup priority, the PCE determines whether a new tunnel can be established by preempting an existing tunnel. The PCE can preempt the existing tunnel if: • Setup priority of the new tunnel is higher than that of the existing tunnel. • Preemption releases enough bandwidth for the new tunnel.

Table 130: Fields on the Add Tunnel Page (*Continued*)

Field	Description
Routing Method	From the list, select a routing method to specify whether the PCE should compute and provision the path for the tunnel: The available options are: • routeByDevice—This is the default routing method when the PCE learns or creates a PCC-controlled tunnel. For this method, The PCE does not compute and provision a path. This method is appropriate for three types of tunnels: RSVP TE PCC-controlled tunnels, segment routing PCEP-based tunnels, and segment routing NETCONF-based tunnels. Use routeByDevice routing method to avoid encountering limitation related to the maximum SID depth (MSD). SR LSP is provisioned with only destination node SID, and additional SID from Strict or Loose hop definition. NOTE: If you select this routing method, your router must run Junos OS Release 19.1 or later. This is to ensure that the router can abide by the hop requirements that you specify in the Path tab in this configuration. • Other routing methods (default, ddelay, IS-IS, OSPF)—When a PCC-controlled tunnel uses a routing method other than RouteByDevice, the PCE computes and provisions the path as a strict explicit route. The tunnel's existing explicit route might be modified to a PCE-computed strict explicit route. For example, a loose explicit route specified by you or learned from the router might be modified to a strict explicit route.
Hold	Specify the hold priority for the tunnel traffic. Priority levels range from 0 (highest priority) through 7 (lowest priority). The default is 7, which is the standard MPLS tunnel definition in Junos OS. Based on the hold priority, the PCE determines whether the tunnel can be preempted or not. If the hold priority for a tunnel is higher, it is unlikely for the tunnel to be preempted.

Table 130: Fields on the Add Tunnel Page (*Continued*)

Field	Description
Color Community	NOTE: This field is available only for the SR provision type. Assign a color for a segment routing tunnel that can be used to map traffic on the tunnel. Range: 0 through 4294967295. A segment routing (SR) policy uses the color community value to steer traffic flows on to a tunnel. By using Color Community, you can steer traffic into SR LSPs.
FlexAlgo ID	Select the flex algo ID that you want to associate with the LSP. This field lists all the flex algo IDs that are associated with both ingress and egress devices in your network. When the LSP is provisioned, the LSP adheres to the pre-defined FAD constraints such as metric type, administrative group restrictions, optimizing traffic paths, and so on. You can select the flex algo ID only for SR (segment routing) and SRv6 LSPs.
CONSTRAINTS	
Admin Group Include All	From the list, select one or more admin group bits for the tunnel to traverse links that include all of the admin groups specified in this field. The maximum selections allowed is 32.
Admin Group Include Any	From the list, select one or more admin group bits. The tunnel traverses links that include at least one of the admin groups specified in this field. The maximum selections allowed is 32.
Admin Group Exclude	From the list, select one or more admin group bits. The tunnel traverses links that do not include any of the admin groups specified in this field. The maximum selections allowed is 32.
Maximum Delay	Specify the maximum delay (in milliseconds) for the tunnel, which is used as a constraint for tunnel rerouting.
Maximum Hop	Specify an integer value for the maximum number of hops that the tunnel can traverse.

Table 130: Fields on the Add Tunnel Page (*Continued*)

Field	Description
Maximum Cost	Specify an integer value for the maximum cost to be associated with the tunnel.
ADVANCED	
Count	Specify the number of parallel LSPs to be created between two endpoints. These LSPs share the same design parameters as specified in the Constraints tab.
Delimiter	NOTE: This field is available only when the count value is greater than 1. Specify the delimiter value. You can use alphanumeric characters, underscore, hyphen, period, or colon. We recommend that you do not use greater than (>), less than (<), and forward slash (/). This value is used in the automatic naming of parallel tunnels that share the same design parameters. Routing Director names the tunnel based on the name you enter in the Properties tab followed by the delimiter value and a unique numerical value beginning with 1. Example: myTunnel_1, myTunnel_2, and so on.
Description	Specify a comment or description for the tunnel for your reference.
IP Z	Select the IP address of the ingress device (Device Z).
PATH	
Routing Path Type	From the list, select the type of routing path for the tunnel: • Dynamic—Allows the PCE to compute a path without imposing any path restrictions. • Required—Prevents the PCE from using any other path for this tunnel. If the required path is not viable and available, the tunnel is down and the PCE does not perform computation to look for an alternate path. • Preferred—Instructs the PCE to use this path over any other, as long as it is viable and available. If it is not viable and available, the PCE computes an alternate path.

Table 130: Fields on the Add Tunnel Page (*Continued*)

Field	Description
Add Hop	This option is available only if the routing path type is Preferred or Required. Click the Add (+) icon or click Add Hop. From the list, select an option as the first hop between node A and node Z. In addition, click the toggle button next to this field to specify whether the hop is strict or loose: • If you specify the hop as strict, the tunnel must take a direct path from the previous router to this router. • If you specify the hop as loose, the tunnel can take any path to reach this router; the PCE chooses the best path. To add more hops, click the + icon again. You can add a maximum of 37 hops. NOTE: When specifying a loose hop, you can choose from all links in the network. When specifying a loose hop for a Required path, anycast group SIDs are also available for selection.

Add Diverse Tunnels

You can pair Label Switched Paths (LSPs) so that they do not overlap at any point between their ingress (starting point) and egress (ending point). This is crucial for network resilience, as it ensures that if one path fails, the other remains operational.

There are three types of tunnel diversity:

- **Link Diversity**—Ensures LSPs do not share common links. If one LSP uses a specific link, the other LSP in the pair will use a different link, thus avoiding shared points of failure.
- **SRLG (Shared Risk Link Group) Diversity**—Ensures LSPs avoid links that share the same risk factors.
- **Site Diversity**—Ensures LSPs avoid the same nodes and therefore the same sites. This is the most comprehensive form of diversity, as it includes both link and SRLG diversity.



NOTE: It is not possible to create a diverse group with more than two tunnels. If you attempt to add more tunnels to a diversity group, the additional tunnels may not maintain the expected diversity.

To implement tunnel diversity in your network:

1. Navigate to **Observability > Network > Topology**.
The Topology page is displayed with the topology map at the center and the network information table at the bottom of the page.
2. Navigate to the Tunnels tab.
3. Click the **Provisioning** drop down and select **Diverse Tunnels**.
The Add Diverse Tunnels page appears.
4. Complete the configuration on each tab according to the guidelines in [Table 131 on page 543](#).



NOTE: Fields marked with an asterisk (*) are mandatory.

5. (Optional) From any tab, click **Preview Path** at the bottom of the page to view the path on the topology map.
6. Click **Add** to add the tunnel.
A confirmation message appears on the top of the page, indicating that diverse tunnels are created.

Table 131: Fields on the Add Tunnel Page

Field	Description
PROPERTIES	

Table 131: Fields on the Add Tunnel Page (*Continued*)

Field	Description
Provisioning Method	From the list, select one of the following methods to be used to provision the tunnel: • PCEP (Path Computation Element Protocol)—The path computation element (PCE) initiates the tunnel and the associated configuration statements do not appear in the router configuration file. Upon provisioning, this tunnel is added as a PCE-initiated tunnel. This is the default provisioning method. • NETCONF—The tunnel is statically provisioned and the associated configuration statements appear in the router configuration file. Upon provisioning, this tunnel is added as a device-controlled tunnel. NOTE: • For Cisco IOS-XR routers, NETCONF-based tunnel provisioning has the same capabilities as PCEP-based tunnel provisioning. • When provisioning tunnels by using NETCONF one at a time, the provisioning order might be sent before the response to a previous provisioning order is received. The second order might not have the correct bandwidth allocation information and the PCE might not be able to provide ECMP. We recommend provisioning multiple tunnels through NETCONF in one operation (bulk provisioning) in order to avoid this issue.
Provision Type	From the list, select the type of tunnel that you want to provision: • RSVP • SR (segment routing) • SRv6
Diversity Group	Specify the name of the diversity group to which this tunnel belongs. If you want to create a diverse path for a pair of tunnels, then ensure that the tunnels belong to the same diversity group. Use alphabets, numbers, period(.), hyphen (-), colon (:), or underscore (_).

Table 131: Fields on the Add Tunnel Page (*Continued*)

Field	Description
Diversity Level	Select one of the following diversity levels : • Link—Two paths don't intersect at any given link. Link diversity is the weakest. • SRLG (Shared Risk Link Group)—Two paths don't intersect at any of the group's links or nodes (aside from the source and destination). SRLG diversity includes link diversity. • Site—Two paths don't intersect at any given site (aside from the source and destination). Site diversity is the strongest as it includes SRLG and link diversity. The diversity level determines the constraints that must be applied during path computation to ensure diverse and optimal paths.
TUNNEL (Specify the following for Tunnel 1 and Tunnel 2.)	
Name	Specify a unique name for the tunnel. You can use any number of alphanumeric characters, period, hyphens, colons, and underscores. NOTE: If you are adding multiple parallel tunnels that will share the same design parameters, the name you specify here is used as the base for automatically naming those tunnels. See the Count and Delimiter fields in the Advanced tab for more information.
Device A	From the list, select the device that you want to use as the ingress node.
Device Z	From the list, select the device that you want to use as the egress node.

Table 131: Fields on the Add Tunnel Page (*Continued*)

Field	Description
Planned Bandwidth	Specify the planned bandwidth (along with valid units, with no space between the bandwidth and units) for the tunnel. If you specify a value without units, bps is automatically applied. Valid units are: • B or b • M or m • K or k • G or g Examples: 50M, 1000b, 25g.
Planned Metric	Specify the static tunnel metric. The PCE uses this metric to route the tunnel instead of allowing the router to choose a path.
Setup	Specify the setup priority for the tunnel traffic. Priority levels range from 0 (highest priority) through 7 (lowest priority). The default is 7, which is the standard MPLS tunnel definition in Junos OS. Based on the setup priority, the PCE determines whether a new tunnel can be established by preempting an existing tunnel. The PCE can preempt the existing tunnel if: • Setup priority of the new tunnel is higher than that of the existing tunnel. • Preemption releases enough bandwidth for the new tunnel.
Hold	Specify the hold priority for the tunnel traffic. Priority levels range from 0 (highest priority) through 7 (lowest priority). The default is 7, which is the standard MPLS tunnel definition in Junos OS. Based on the hold priority, the PCE determines whether the tunnel can be preempted or not. If the hold priority for a tunnel is higher, it is unlikely for the tunnel to be preempted.

Table 131: Fields on the Add Tunnel Page (*Continued*)

Field	Description
Color Community	NOTE: This field is available only for the SR provision type. Assign a color for a segment routing tunnel that can be used to map traffic on the tunnel. Range: 0 through 4294967295. A segment routing (SR) policy uses the color community value to steer traffic flows on to a tunnel. By using Color Community, you can steer traffic into SR LSPs.
FlexAlgo ID	Select the flex algo ID that you want to associate with the LSP. This field lists all the flex algo IDs in your network. When the LSP is provisioned, the LSP adheres to the pre-defined FAD constraints such as administrative group restrictions, optimizing traffic paths, and so on. You can select the flex algo ID only for SR (segment routing) and SRv6 LSPs.
CONSTRAINTS	
Admin Group Include All	From the list, select one or more admin group bits for the tunnel to traverse links that include all of the admin groups specified in this field. The maximum selections allowed is 32.
Admin Group Include Any	From the list, select one or more admin group bits. The tunnel traverses links that include at least one of the admin groups specified in this field. The maximum selections allowed is 32.
Admin Group Exclude	From the list, select one or more admin group bits. The tunnel traverses links that do not include any of the admin groups specified in this field. The maximum selections allowed is 32.
Maximum Delay	Specify the maximum delay (in milliseconds) for the tunnel, which is used as a constraint for tunnel rerouting.
Maximum Hop	Specify an integer value for the maximum number of hops that the tunnel can traverse.

Table 131: Fields on the Add Tunnel Page *(Continued)*

Field	Description
Maximum Cost	Specify an integer value for the maximum cost to be associated with the tunnel.

About the FlexAlgo Tab

IN THIS SECTION

- [Tasks You Can Perform | 548](#)

To access this tab, click **Observability > Network > Topology > FlexAlgo**.

Routing Director discovers FADs that you have configured on the devices in your network.

On this tab, you can view the details of the Flexible Algorithm Definition (FAD). If you select a flex algo ID in the flex algo table, then the nodes and links that are participating in the selected flex algo ID is highlighted on the topology map.



NOTE: Before you configure flex algo on the device, we recommend that you onboard the device using Routing Director and ensure that BGP-LS is configured.

For information about how to configure flex algo on the devices in your network, see [How to Configure Flexible Algorithms in OSPF for Segment Routing Traffic Engineering](#) and [Flexible Algorithms in IS-IS for Segment Routing Traffic Engineering](#).

[Table 132 on page 549](#) provides information on the fields in the FlexAlgo tab.

Tasks You Can Perform

- View details about a FlexAlgo ID—To view details (such as source node, metric type, and so on):
 - Hover over a FlexAlgo ID row in the table and click the **Details** icon. On the **FlexAlgo ID-Number** page that appears, you can view the details of the selected flex algo.

- Select the FlexAlgo ID on the table and click **More > Show Detail**. On the **FlexAlgo ID-Number** page that appears, you can view the details of the selected flex algo.
- View the computed path on the Topology map—When you select a flex algo ID, the links and nodes that are participating in the selected flex algo ID are highlighted on the Topology map. The color conventions are as follows:
 - Red highlight—Indicates one-sided participation from the device for the selected FlexAlgo ID.
 - Yellow line—Indicates both the devices are participating for the selected the FlexAlgo ID.
- View conflicts in FAD parameters—Sometimes, multiple devices in your network might participate in the same FlexAlgo definition (FAD). Even though the FAD ID is the same, the user might have configured FAD parameters differently at the device level. In such scenarios, the FlexAlgo table highlights the differences in FAD parameters in red so that you can easily identify the difference in FAD parameters configuration.



NOTE: The conflicts are highlighted only for the following FAD parameters: Metric type and calculation type.

- To perform common actions such as filtering using advanced filter criteria and resetting preferences, see ["Network Table Overview" on page 515](#).

Table 132: Fields on the FlexAlgo Table

Fields	Description
FlexAlgo ID	Displays the FlexAlgo definition (FAD) identifier number, which ranges from 128 through 255.
Source Node	Displays the device name on which the flex algo is configured.
Metric Type	Displays the metric type (IGP, TE, or Delay) that you have defined in the flex algo definition.

Table 132: Fields on the FlexAlgo Table *(Continued)*

Fields	Description
Calc Type	Displays the calculation type that you have defined in the flex algo definition. • SPF—Influences the SPF computation on your device based on a certain local policy such as traffic engineering shortcuts. • Strict SPF—Local policy cannot influence the SPF path selection.
Priority	Displays the priority that you have defined in the flex algo. By default, the priority is set as 0.
Admin Group Exclude	Lists the administrative groups that are excluded.
Admin Group Include Any	Includes any of the administrative groups that listed.
Admin Group Include All	Includes all of the administrative groups that listed.
Redundancy Check	If the same FlexAlgo definition (FAD) ID is applicable for at least two different devices then the redundancy check is displayed as Yes. If the FAD ID is set for only one of the devices then the redundancy check is displayed as No.

RELATED DOCUMENTATION

[Device Onboarding Workflow | 173](#)

[Configuring Topology Acquisition Using BGP-LS | 559](#)

Reroute LSPs

Routing Director uses Path Computation Element Protocol (PCEP) or Network Configuration Protocol (NETCONF) to learn and discover Label-switched Paths (LSPs) in your network.

Routing Director can automatically reroute LSP based on the following predefined parameters that you set at your organization level:

- **Equal-cost multipath (ECMP) placement method**—Routing Director adopts ECMP-based rerouting of LSPs to load balance traffic and increase bandwidth by fully utilizing unused bandwidth on links to the same destination.

You need to set the **ECMP Placement Method** parameter. Based on the placement method you select (Random, Least-Filled, and Most-Filled), Routing Director evaluates link metrics, and if there are multiple ECMP path available towards the same destination then Routing Director selects a path based on the user settings.

- **Link utilization threshold**—Routing Director automatically reroutes LSP when the utilization of a link exceeds a specified threshold, ensuring optimal network performance and preventing congestion. For automatic rerouting of LSPs, you need to set **Link Utilization Threshold** and **Threshold Rerouting Interval** parameters.

For example, if the link utilization threshold is set at 60%, Routing Director immediately reroutes traffic if the link utilization exceeds this value. Even after the rerouting, if the link utilization is more than 60% , then Routing Director waits for the duration you have specified in Threshold Rerouting Interval (e.g., 3 minutes) before rerouting to avoid constant adjustments for transient traffic spikes.



NOTE: Besides organization-level settings, you can also set the link utilization threshold at the link level. The link-level setting overrides the organization level settings.

- **Path optimization interval**—Routing Director ensures LSPs are routed through the most optimal paths. Even if an LSP is rerouted due to threshold crossing, link failure or topology change, the LSP should return to its optimal path once the congestion is resolved. For Routing Director to reroute LSPs through the most optimal paths, you need to set the **Path Optimization Interval** parameter.

To configure the parameters at your organization level:

1. Click **Settings Menu > System Settings** on the banner.

The Organization Settings page appears.

2. Configure or modify the following pathfinder settings:

Parameter	Description
ECMP Placement Method	Select a placement method for equal-cost multipath (ECMP) routing. • Random—Randomly selects one of the ECMP paths

Parameter	Description
	- Least filled—Selects the path with the maximum available bandwidth - Most filled—Selects the path with the minimum available bandwidth Based on the placement method you opt, Routing Director evaluates link metrics, distributes traffic, and improves bandwidth utilization on multiple links to the same destination.
Path Optimization Interval	Specify the interval after which Routing Director must automatically trigger path optimization. To disable automatic path optimization, enter 0 second. Range: 0 through 9 minutes (m), hours (h) or days (d).
Link Utilization Threshold	Specify the threshold value (in percentage) for link utilization. When traffic on a link exceeds this value, Routing Director reroutes label-switched path (LSPs). If you do not specify a threshold, the LSP is not rerouted. If you specify 0, the links are blocked. Range: 0 through 100 percentage.
Threshold Rerouting Interval	Specify the minimum interval (in minutes) after which Routing Director reacts to any traffic or delay violations. If you don't specify an interval, the label-switched path (LSP) is not rerouted in case of a violation. Range: 1 through 300 minutes

3. Click **Save** to save the settings.

Verify that the settings are saved and close the Organization Settings page.

Add and Remove LSP Delegation

IN THIS SECTION

- [Add LSP Delegation | 554](#)
- [Remove LSP Delegation | 554](#)

You can use the Configure LSP Delegation page (**Observability > Network > Topology > Tunnel tab > Delegation > Configure Delegation**) to delegate the management of Path Computation Client-controlled label-switched paths (PCC-controlled LSPs) to the Path Computation Element (PCE). After delegation, these LSPs are managed by the PCE and are classified as PCC-delegated LSPs. You can also return the control of these LSPs to the PCC by removing the delegation, and these LSPs are then re-classified as PCC-controlled LSPs.

Understand the Behavior of Delegated and Undelegated LSPs

In both standalone (where the Routing Director components run on a single primary node) and high availability (HA) cluster configurations (where the Routing Director components run on multiple primary nodes), whenever a Path Computation Element Protocol (PCEP) session on a PCC goes down, all the LSPs that originated from that PCC are removed from the database, except those LSPs with parameters configured from the GUI.

For PCC-controlled LSPs, all the configuration changes made to the LSP attributes from the router (that is, the PCC) are stored only in the router's configuration file. When you delegate such LSPs to the PCE and change the LSP attributes from the GUI, the changes are stored only in the database. The PCE communicates the changes to the PCC by using PCEP.

You can still make changes to the LSP attributes from the router and these changes are stored in the router's configuration file, but the changes do not take effect as long as the LSPs remain delegated to the PCE.

The configuration changes made from the PCC take effect automatically only after the LSPs are undelegated from the PCE, which means that the control of the LSPs is returned to the PCC. After undelegation, the PCE obtains the LSP state from the PCEP report messages. If you change the LSP attributes from the router and the PCE isn't available at this time, the PCE obtains the latest LSP state after it becomes available. Therefore, if you delegate such LSPs to the PCE, the LSPs carry the latest state.

If you perform the Undelegate operation on a delegated LSP, the Path Computation Server (PCS) uses the bandwidth reported by the device for the planned bandwidth instead of the user input value.

Add LSP Delegation

To delegate one or more LSPs to the PCE:

1. Navigate to **Observability > Network > Topology > Tunnel tab > Delegation > Configure Delegation**.
The Configure Delegation page is displayed.
2. From the **Add Delegation** tab, select the LSPs that you want to delegate to the PCE.
To select all the available LSPs for delegation, click the check-box beside the **Name** column.
3. Click **Delegate LSP**.

A confirmation message appears on the top of the page indicating that the LSP delegation was requested.

The `lsp-external-controller pccd` statement is added to the router configuration and the LSPs are delegated to the PCE.

The Control Type column in the Tunnel tab (**Topology > Tunnel**) of the network information table displays **Delegated** for the selected LSPs.



NOTE:

- If there is a failure along the path, the PCE reroutes the delegated LSPs (RSVP LSPs and SR LSPs with a single primary path) around the failed network element.
- When the PCE doesn't find a path for the delegated SR LSP, the PCE sends a PCEP message to the PCC with an empty explicit route object (ERO). The PCC then sets the LSP's operational status as Down and sends a PCEP message to the PCE indicating that the LSP's operational status is Down. Later, when the PCE computes a path for the LSP, the PCE sends a PCEP message to the PCC, with the ERO. The PCC then provisions the ERO and sends a PCEP message to the PCE indicating that the LSP's operational status is now Up.

The empty ERO feature for delegated SR LSPs and delegated RSVP LSPs is supported only for Juniper Networks PCC devices and this feature has not been tested with other PCC devices. For delegated SR LSPs, this feature is supported only on Juniper Networks PCC devices running Junos OS Release 20.4R1 or later.

Remove LSP Delegation

To undelegate one or more LSPs from the PCE:

1. Navigate to **Observability > Network > Topology > Tunnel tab > Delegation > Configure Delegation**.
The Configure Delegation page is displayed.
2. From the **Remove Delegation** tab, select the LSPs that you want to undelegate.
To select all the available LSPs for undelegation, click the check-box beside the **Name** column.

3. Click **Undelegate LSP**.

A confirmation message appears on the top of the page indicating that the LSP delegation change was requested.

The delegation statement is removed from the router configuration and the control of these LSPs is returned to the PCC.

The Control Type column in the Tunnel tab (**Topology > Tunnel**) of the network information table displays **Device Controlled** for the selected LSPs.

Dynamic Topology Workflow

This workflow describes the tasks that you need to perform for the dynamic topology feature.

The dynamic topology feature enables you to view the dynamic changes (real-time) in your topology even if the devices are not managed by Juniper Routing Director. You can view the real-time changes on the Topology page (**Observability > Network > Topology**).

To view real-time changes in your network, do the following:

1. Establish BGP-LS peering.

- a. Configure BGP-LS on the routers. For more information, see ["Configuring Topology Acquisition Using BGP-LS" on page 559](#).
- b. Login to the Juniper Routing Director GUI, and do the following:
 - i. Navigate to the Topology page (**Observability > Network > Topology**).
 - ii. On the Topology Menu bar, a vertical bar at the top-right corner of the Topology page, click the Settings icon.
 - iii. Specify the IP address of the BGP-LS peer and autonomous system (AS) number so that Juniper Routing Director can establish a BGP-LS peering session.

2. Establish PCEP sessions.

- a. Specify a VIP address either at the time of installing Routing Director or post-installation of the cluster. For more information, see ["Configure a PCE Server" on page 556](#)
- b. Configure PCEP on each PE router. For more information, see ["Configuring PCEP on a PE Router" on page 557](#).

You can now view the real-time changes in your network on the Topology page.

Configure a PCE Server

Routing Director uses Path Computational Element Protocol (PCEP) sessions to discover tunnels that run in your network. To establish PCEP sessions between Routing Director and the devices managed by it, you must configure a VIP address using Paragon Shell. The VIP address can be configured either during installation of the Routing Director cluster or post-installation of the cluster.

To configure the PCE Server VIP address using Paragon Shell:

1. Log in to installer node of the Routing Director cluster.
2. Type configure to enter configuration mode.
3. Configure the PCE Server VIP address, commit, and exit the configuration mode.

```
root@Primary1# set paragon cluster applications pathfinder pce-server pce-server-vip pce-server-vip

[edit]
root@Primary1# commit
commit complete

[edit]
root@Primary1# exit
Exiting configuration mode

root@Primary1>
```

Where:

pce-server-vip is the VIP address used by the PCE server.

4. Regenerate the configuration files.

```
root@Primary1> request paragon config
Paragon inventory file saved at /epic/config/inventory
Paragon config file saved at /epic/config/config.yml
```

5. Redeploy the PCE server application.

```
root@Primary1> request paragon deploy cluster input "-t metallb,addon-apps -e
target_components=pceserver"
Process running with PID: 23xx022
To track progress, run 'monitor start /epic/config/log'
```

After you configure the PCE server, you can view tunnel-related information on the Tunnels tab of the network table (**Observability > Network > Topology > Tunnels**). For more information, see ["About the Tunnels Tab" on page 528](#).

RELATED DOCUMENTATION

| *Post Installation Tasks*

Configuring PCEP on a PE Router

IN THIS SECTION

- [Configuring a PE Router as a PCC | 557](#)

A Path Computation Client (PCC) supports the configurations related to the Path Computation Element (PCE) and communicates with Routing Director. However, you must configure PCEP on each PE router to configure the router as a PCC and establish a connection between the PCC and Routing Director. A PCC initiates path computation requests, which are then executed by Routing Director.

Configuring a PE Router as a PCC

To configure a PE router as a PCC:

1. Enable external control of LSPs from the PCC router to Routing Director.

```
[edit protocols]
user@PE1# set mpls lsp-external-controller pccd
```

2. Specify Routing Director (**paragon**) as the PCE that the PCC connects to, and specify the Routing Director host external IP address as the destination address.



NOTE: The IP address that you specify here is the VIP address (*pce-server-vip*) that you have configured either during installation of the Routing Director cluster or post-installation of the cluster.

```
[edit protocols]
user@PE1# set pcep pce paragon destination-ipv4-address 10.99.99.1
```

3. Configure the destination port for the PCC router that connects to Routing Director (PCE server) using the TCP-based PCEP.

```
[edit protocols]
user@PE1# set pcep pce paragon destination-port 4189
```

4. Configure the PCE type.

```
[edit protocols]
user@PE1# set pcep pce paragon pce-type active
user@PE1# set pcep pce paragon pce-type stateful
```

5. To verify that PCEP has been configured on the router, open a telnet session to access the router, and run the following commands:

```
user@PE1> show configuration protocols mpls
```

Sample output:

```
lsp-external-controller pccd;
```

```
user@PE1> show configuration protocols pcep
```

Sample output:

```
pce paragon {
  local-address 10.0.0.101;
  destination-ipv4-address 10.99.99.1;
```

```

destination-port 4189;
pce-type active-stateful;
lsp-provisioning;
}

```

Configuring Topology Acquisition Using BGP-LS

To configure BGP-LS topology acquisition on Juniper routers, perform the following configuration steps:

1. Specify the autonomous system (AS) number for the node (BGP peer).

```

[edit routing-options]
user@paragon_junosvm# set autonomous-system AS_number

```

2. Specify the BGP group name and type for the node.

```

[edit protocols bgp]
user@paragon_junosvm# set group group_1 type internal

```

3. Specify a description for the BGP group for the node.

```

[edit protocols bgp group group_1]
user@paragon_junosvm# set description "Paragon BGP-TE Peering"

```

4. Enable the traffic engineering features for the BGP routing protocol.

```

[edit protocols bgp group group_1]
user@paragon_junosvm# set family traffic-engineering unicast

```

5. Specify the IP address for the installer node from which Routing Director is deployed.

```

[edit protocols bgp group group_1]
user@paragon_junosvm# set neighbor paragon-automation-node-ip-address

```


Enable PCEP Security

You can enforce security between the PCE (Path Computation Element) server and Path Computation Clients (PCC) by configuring a security mode for the PCEP connections. Routing Director supports the following security modes:

- Auto Detect—Accepts PCEP connection regardless of whether you have security enabled or not.
- Strict Disable—Accepts only non-secured PCEP connections.
- Strict Enable—Enforces security on all PCEP sessions. If a device does not have security enabled, its connection is rejected.

To enforce security between the PCE server and PCC using Paragon Shell:

1. Log in to a cluster node. You are logged in to Paragon Shell.
2. Type `configure` to enter configuration mode.

```
user@node>
user@node> configure
Entering configuration mode
[edit]
user@node#
```

3. By default, the security mode is *strict-disable*. Set the **pce_server_global_default_tls_mode** parameter to *auto-detect* or *strict-enable* so that the PCEP security is enabled.

```
user@node# set paragon cluster applications pathfinder pce-server pce-server-global-default-
tls-mode strict-enable
```



NOTE: You can update the **pce_server_global_default_tls_mode** parameter before or after installing Routing Director.

4. If you enable secured connections (*auto-detect* or *strict-enable*), you can use either system-generated certificates or custom certificates.

If you are using system-generated certificate:

- The server certificate (**tls.crt**) and private key (**tls.key**) are automatically generated. You can view the system-generated certificate using the following command:

```
user@node# kubectl get secrets | grep pce-default-cert
```



NOTE: System-generated certificates are managed by cert-manager, which automatically renews certificates 90 days before their expiration. The PCE server seamlessly detects the updated certificate, ensuring that existing PCEP sessions remain uninterrupted while continuing to operate securely.

If you are using custom certificates for added security, use the following steps to upload a custom certificate:

- a. Set the **use_custom_certificate** parameter to *true*.

```
user@node# set paragon cluster applications pathfinder use_custom_certificate true
```

- b. Upload the custom certificate and the private key in the **/root/epic/config** folder of the node from which the cluster is deployed.

```
user@node# set paragon cluster applications pathfinder pce-server pce-server-user-
certificate pce-server-custom-certificate-filename certificate_filename
user@node# set paragon cluster applications pathfinder pce-server pce-server-user-
certificate pce-server-custom-certificate-key-filename certificate_private_key
```

- c. You can view the custom certificate using the following command:

```
user@node# kubectl get secrets | grep pce-custom-cert
```



NOTE: Custom certificates are not automatically renewed. You need to upload the renewed server certificate and private key at least a few days before the expiry of the certificate.

5. Commit and quit configuration mode.

```
user@node# commit and-quit
commit complete
Exiting configuration mode
user@node>
```

6. Copy the certificate on the device using the following procedure:

- a. Type exit from Paragon Shell to exit to the Linux root shell.
- b. If you are using a system-generated certificate, download the certificate using the following command:

```
user@node:~#kubectl get secret -n northstar pce-default-cert -o jsonpath='{.data.ca\.crt}'
| base64 -d > ns1.pem
```

- c. Upload the certificate (system-generated or custom) to the device using the following command:

```
user@node:~#scp ns1.pem root@device-IP-address:/tmp
```

- d. Log in to the device to configure and commit the following:

```
user@device# set security pki ca-profile pccd-tls ca-identity pccd-tls
user@device# set security pki ca-profile pccd-tls revocation-check disable
user@device# commit and-quit
```

- e. Load the copied certificate file.

```
user@device> request security pki ca-certificate load ca-profile pccd-tls filename /tmp/
ns1.pem
Fingerprint:
    34:3a:cf:3e:9e:5a:fb:89:e5:53:15:75:79:b1:31:6a:e3:ed:77:55 (sha1)
    2a:23:48:70:b7:ad:db:c7:5b:dc:58:e5:16:71:aa:da (md5)
    :83:66:f1:6e:55:15:59:ad:a9:60:07:0a:08:0a:f7:f2:2f:90:53:1a:65:81:ed:09:87:2c:ce:3b:31
:d5:f9 (sha256)
Do you want to load this CA certificate ? [yes,no] (no) yes
```

- f. Configure the PCEP mode and commit the configuration.

```
user@device# set protocols pcep tls-strict
user@device# commit and-quit
```

- g. Delete certificate authority (CA) digital certificates from the device.

```
user@device> clear security pki ca-certificate ca-profile pccd-tls
```

7. Do one of the following:

- a. If you have not installed Routing director, run the following commands to configure the PCE server based on the **pce_server_global_default_tls_mode** setting:

```
user@node> request paragon config
user@node> request paragon deploy cluster
```

- b. If you have already installed Routing director, run the following commands to redeploy only the PCE server component:

```
user@node> request paragon config
user@node> request paragon deploy cluster input "-t addon-apps -e
target_components=pceserver"
```

The PCEP Status column on the Device tab of the Topology page (**Observability > Network**) is displayed as *Up* if you have enabled PCEP security and have uploaded certificates on both the server and the client.

If you encounter issues, such as a PCEP session is down, ensure that you:

- Check the logs on both the PCE server and the client device to identify any errors related to the certificate or security configuration.
- Ensure that the correct mode is set (Auto Detect or Strict Enable) and that the required certificates are properly configured.
- Verify that the certificates on both the server and the client match and are valid.
- If you have performed a disaster recovery, ensure that you enable PCEP security again.

RELATED DOCUMENTATION

| [set paragon cluster applications](#) | 1498

Routing Observability

IN THIS CHAPTER

- [Routing Observability Overview | 565](#)
- [Enable Routing Observability and View Routing Analytics Data | 566](#)
- [About the Route Explorer Page | 566](#)
- [About the Routing Status Tab | 567](#)
- [About the Routing Updates Tab | 571](#)
- [About the Adjacencies Tab \(Route Explorer\) | 573](#)

Routing Observability Overview

The Routing Observability feature in Routing Director allows you to actively monitor the overall routing health of your network in real time. Routing Director provides a Route Explorer dashboard that you can use to monitor routing health. However, to be able to use this feature, you must:

- Configure a VIP (by using Paragon Shell) which is used by external network devices as BMP station IP address to establish the BMP session. See *Deploy the Cluster*.
- Enable routing protocol analytics, and configure BGP data collection when you add a device profile. See ["Add a Device Profile" on page 208](#).

Route Explorer

The Route Explorer dashboard allows you to monitor all routing prefixes in the network. You can track historical route events and route attribute details, enabling identification and troubleshooting of reachability issues and route fluctuations. You can also monitor, in real time, BGP peers and routing information base (RIB) statistics, allowing you to detect and resolve routing control plane issues. You can use the Route Explorer dashboard to consistently monitor and analyze peer information and understand their effect on the networks routing performance. See ["About the Route Explorer Page" on page 566](#).

Benefits

- Monitor the overall routing health of your network by using the Route Explorer dashboard, which presents real-time routing data
- Identify routing issues and take corrective actions wherever necessary

Enable Routing Observability and View Routing Analytics Data

The Routing Observability feature in Routing Director allows you to actively monitor the overall routing health in real time. You can view this data on the Route Explorer dashboard of the Routing Director GUI.

Perform these steps to enable the routing observability feature and view routing analytics data:

1. Configure a virtual IP (VIP) address when you configure the Routing Director cluster to enable the routing observability feature.

Configure a VIP (by using Paragon Shell) which is used by external network devices as the BMP station IP address to establish the BMP session. For more information, see *Deploy the Cluster*

2. Enable routing protocol analytics on the Create Device Profile page when you create a device profile.

Enable the **Routing Protocol Analytics** option when you add a device profile to configure routing data collection. For more information, see ["Add a Device Profile" on page 208](#).

3. Configure routing data collection.

Configure BGP data collection parameters such as BGP state and routing information bases (RIBs) to view BGP analytics data on the Route Explorer dashboard. For more information, see ["Add a Device Profile" on page 208](#).

4. View and monitor routing analytics.

Use the Route Explorer dashboard to monitor the health and status of IPv4 and IPv6 routes of all BGP speakers. For more information, see ["About the Route Explorer Page" on page 566](#).

About the Route Explorer Page

To access this page, click **Observability > Routing > Route Explorer**.

Users with Super User, Network Admin, and Observer roles can monitor the overall routing health of the network in real time by using the Route Explorer dashboard.



IMPORTANT: To use the Route Explorer dashboard and its features, you must enable **Routing Protocol Analytics** and configure BGP data collection when you create a device profile. See ["Add a Device Profile" on page 208](#) for more information.

This dashboard consists of the following tabs:

- ["Routing Status" on page 567](#)—View the states of all routing prefixes and peers.
- ["Routing Updates" on page 571](#)—View a record of all routes added to and deleted from the RIB.
- ["Adjacencies \(Route Explorer\)" on page 573](#)—View BGP peers in the network for which BGP state and BGP routing information base (RIB) monitoring is enabled.

You can also view a time series graph of the total number of routes, and routing events. These graphs can be viewed for a predefined or a custom period.

About the Routing Status Tab

IN THIS SECTION

- [About the Prefixes Tab | 567](#)
- [About the Devices Tab | 568](#)
- [Tasks You Can Perform | 569](#)

To access this tab, click **Observability > Routing > Route Explorer > Routing Status**.

Use this tab to view details of all routing prefixes in the routing information base (RIB) and BGP speakers in the network. You can view this information only when BGP state and RIB monitoring is enabled when you create a device profile. For more information, see ["Add a Device Profile" on page 208](#). The Routing Status tab consists of the ["Routing Status > Prefixes" on page 567](#) and ["Routing Status > Devices" on page 568](#) tabs.

About the Prefixes Tab

To access this tab, click **Observability > Routing > Route Explorer > Routing Status > Prefixes**.

You can view details of all routing prefixes in the RIB, and the corresponding BGP speakers in the network on this tab.

See ["Tasks You Can Perform" on page 569](#) for more information about the tasks that you can perform from this tab.

Fields Descriptions

[Table 133 on page 568](#) lists the fields on the Routing Status > Prefixes tab.

Table 133: Fields on the Routing Status > Prefixes Tab

Field	Description
Prefix Address	Network prefix address.
Prefix Length	Length of the prefix.
Device Count	Total number of RIB-monitored devices where this prefix is installed as an active route.
AF Type	The address family (AF) type associated with the prefix. The AF types are IPv4-Unicast and IPv6-Unicast.

About the Devices Tab

To access this tab, click **Observability > Routing > Route Explorer > Routing Status > Devices**.

You can view details of all BGP speakers in the network for which BGP state and/or RIB monitoring is enabled. You can also view route add and delete events, and states of all BGP speakers.

See ["Tasks You Can Perform" on page 569](#) for more information about the tasks that you can perform from this tab.

Field Descriptions

[Table 134 on page 569](#) lists the fields on the Routing Status > Devices tab.

Table 134: Fields on the Routing Status > Devices Tab

Field	Description
Hostname	Host name of the device
Current Peers	Total number of BGP peers on the device.
Current Active Routes	Total number of active local routing information base (loc-RIB) routes on the device.
Monitoring	Type of BGP monitoring used. The types are: State (BMP Peer Messages), RIB (BMP RIB Messages)
Status	Current status of the device's BMP session.
Site	Site (geographical location) where the device is installed.

Tasks You Can Perform

This section provides information on the tasks that you can perform from the **Routing Status > Prefixes** and **Routing Status > Devices** tabs.

Common tasks

In the time series graph above the Routing Status table, you can view a graphical representation of:

- The total count of all IPv4 or IPv6 routes of all RIB-monitored devices in the network.

You can also:

- Select the option buttons provided in the top-center of the graph to switch between IPv4 (default) and IPv6 prefixes.
- Click the **Linear** or **Log** buttons in the top-right corner of the graph to switch between a linear or log graphical view.
- View this information for a predefined (2h, 4h, 6h, 8h, 10h, 12h) or a custom time period.

- Hover over any point on the graph to view the state of the prefix at that point.

Tasks you can perform from the Routing Status > Prefixes table

From the Routing Status > Prefixes table, you can:

- View details of all prefixes available.
- Sort, resize, or re-arrange columns in a table (grid).
- Show or hide columns in the table or reset page preferences, using the vertical ellipsis menu.
- Filter the data displayed in the table—Click the filter icon (funnel) and select whether you want to show or hide advanced filters. You can then add or remove filter criteria, save criteria as a filter, apply or clear filters, and so on. The filtered results are displayed on the same page.

You can filter by prefix length, prefix address, and AF type data.

- Click the **Download** button to download all device routing information as a CSV file.
- Click the prefix hyperlink to view a graph of the events (such as routes added and routes deleted) related to that prefix for the past 12 hours, on the Routing Updates tab.

The Routing Updates table is updated with details of the prefix that you selected.

Tasks you can perform from the Routing Status > Devices table

From the Routing Status > Devices table, you can:

- View details and BGP Monitoring Protocol (BMP) status of all devices available.
- Sort, resize, or re-arrange columns in a table (grid).
- Show or hide columns in the table or reset page preferences, using the vertical ellipsis menu.
- Filter the data displayed in the table—Click the filter icon (funnel) and select whether you want to show or hide advanced filters. You can then add or remove filter criteria, save criteria as a filter, apply or clear filters, and so on. The filtered results are displayed on the same page.

You can filter by host name, status, and site data.

- Click the host name hyperlink to view a graph of the events (such as routes added and routes deleted) related to that device for the past 12 hours, on the Routing Updates tab.

The Routing Updates table is updated with details of the device that you selected.

About the Routing Updates Tab

IN THIS SECTION

- [Tasks You Can Perform | 571](#)
- [Field Descriptions | 572](#)

To access this tab, click **Observability > Routing > Route Explorer > Routing Updates**.

Use this tab to view a record of all routes add to and deleted from the local routing information base (loc-RIB) of the BGP speakers in the network. You can view this information only for those devices where RIB monitoring is enabled when you create a device profile. For more information, see ["Add a Device Profile" on page 208](#). Route add and delete events that have occurred on the monitored devices within the past 24 hours are displayed in the Route Updates table of this tab. You can also view a graphical representation of all add and delete events.

Tasks You Can Perform

You can perform the following tasks from this tab:

- In the graph above the Routing Updates table, you can view a graphical representation of:
 - All route add and delete events across all monitored devices, by default, that have occurred within the last 12 hours.

You can also view this information for a predefined (2h, 4h, 6h, 8h, 10h, 12h) or a custom time period (of 12-hour intervals).

 - To view a filtered record of route add and delete events (of various devices) that match the prefix that you selected on **Routing Status > Prefixes** tab.
 - To view a filtered record of route add and delete events of a specific device that you selected on **Routing Status > Devices** tab.
- On the Routing Updates table, you can:
 - View a historical record of all routes added to and deleted from the loc-RIB.

You can view this information only for those devices where RIB monitoring is enabled

 - Sort, resize, or re-arrange columns in a table (grid).

- Show or hide columns in the table or reset page preferences, using the vertical ellipsis menu.
- Filter the data displayed in the table—Click the filter icon (funnel) and select whether you want to show or hide advanced filters. You can then add or remove filter criteria, save criteria as a filter, apply or clear filters, and so on. The filtered results are displayed on the same page.

You can filter by reported time, prefix address, prefix length, route distinguisher and hostname data.

Field Descriptions

[Table 135 on page 572](#) lists the fields on the Routing Updates tab.

Table 135: Fields on the Routing Updates Tab

Field	Description
Prefix Address	Network prefix address.
Prefix Length	Length of the prefix.
Local Pref	Local preference of the route.
Hostname	Hostname of the device.
Peer IP	Peer IP address.
Event	Route add or delete events reported on the BGP speaker.
Reported Time	Time the last add or delete event that occurred on the BGP speaker was reported.
Next Hop	Next hop IP address.
Med	Multi-exit discriminator (MED) of the route.

Table 135: Fields on the Routing Updates Tab *(Continued)*

Field	Description
AS Path	Sequence of autonomous system (AS) that the route has passed through to reach the BGP speaker.
AS Origin	Originator autonomous system (AS) number of the BGP peer group.
Community	BGP community that the route is tagged to.
AF Type	Address family (AF) type associated with the prefix.

About the Adjacencies Tab (Route Explorer)

IN THIS SECTION

- [Tasks You Can Perform | 573](#)
- [Field Descriptions | 574](#)

To access this tab, click **Observability > Routing > Route Explorer > Adjacencies**.

Use this tab to view the current state of all BGP peers in the network. You can view this information when either BGP routing information base (RIB) or BGP State monitoring has been enabled when you create a device profile. For more information, see "[Add a Device Profile](#)" on page 208.

Tasks You Can Perform

You can perform the following tasks from this tab:

- View peering information for devices in the network where either BGP RIB or BGP State monitoring is enabled.
- Sort, resize, or re-arrange columns in a table (grid).

- Show or hide columns in the table or reset page preferences, using the vertical ellipsis menu.
- Filter the data displayed in the table—Click the filter icon (funnel) and select whether you want to show or hide advanced filters. You can then add or remove filter criteria, save criteria as a filter, apply or clear filters, and so on. The filtered results are displayed on the same page.

You can filter by hostname and peer IP data.

Field Descriptions

Table 136 on page 574 lists the fields on the Adjacencies tab (Route Explorer).

Table 136: Fields on the Adjacencies Tab (Route Explorer)

Field	Description
Device	Hostname of the device.
Peer IP	Peer IP address.
Peer AS	Autonomous system (AS) number of the BGP peer group. For more information, see Autonomous Systems for BGP Sessions .
State	Current state of the BGP peering session. The most common states are Active , which indicates an attempt in establishing the BGP connection, and Established , which indicates a successful session setup. The other states are transition states such as Connect and Idle , and BGP sessions normally do not stay in those states for extended periods of time.
RIB-In (Pre)	Total number of routes received by the BGP peer before inbound policy is applied.
RIB-Out (Post)	Total number of routes advertised by the BGP peer after inbound policy is applied.

Table 136: Fields on the Adjacencies Tab (Route Explorer) *(Continued)*

Field	Description
Routing Alerts	Details of routing alerts generated.
Down Count	Total number of times the BGP peer has gone down since BGP monitoring was enabled.
Down Reason	Reason for the BGP peer going down.
Reported Time	Time the BGP peer was reported to have gone down.
Local AS	Local autonomous system (AS) number of the BGP peer group. For more information, see Autonomous Systems for BGP Sessions .
Remote Device	IP address of the remote BGP peer.
Local Address	IP address of the local BGP peer.

Rules

IN THIS CHAPTER

- [Rules Overview | 576](#)
- [About the Custom KPI Collection Page | 578](#)
- [About the Rule Instantiations Tab | 578](#)
- [About the Custom Rules Tab | 580](#)
- [About the Alerts Tab \(Custom Rules\) | 584](#)
- [About the Monitor Instance-Name Page | 586](#)
- [Manage Custom Rules | 590](#)
- [Manage Rule Helper Files | 593](#)
- [Add a Rule Instance | 594](#)
- [Manage Rule Instances | 596](#)
- [View KPI Performance and Alerts \(Custom Rules\) | 598](#)

Rules Overview

IN THIS SECTION

- [Understand Custom Rules | 577](#)

A rule is used to extract information from a device in a network. Routing Director enables you to upload custom rules that define how key performance indicators (KPIs) for various device components are collected. You can then apply a custom rule on devices by creating a rule instance. After you deploy a rule instance,

- Alerts are generated when a KPI crosses the threshold limit defined in the custom rules file.

- You can take corrective measures to address the alerts.

Understand Custom Rules

You use a custom rule to define key performance indicators (KPIs). You can organize these rules into rule instances and apply the rule to devices. Before you can upload a custom rule by using the Routing Director GUI, you must define a custom rule. To define a custom rule based on your requirements, we recommend that you contact [Juniper Networks Professional Services](#).

After you have received the custom rule from the Juniper Networks Professional Services team, you must:

1. Upload a custom rule. See ["Upload a Custom Rule" on page 590](#).

2. Include the custom rule in a rule instance.

See step ["3.c" on page 595](#) of the Add a Rule Instance topic.

3. Apply the instance to one or more devices.

See step ["5" on page 595](#) of the Add a Rule Instance topic.

4. Execute or deploy the rule instance.

See step ["8.c" on page 596](#) of the Add a Rule Instance topic.

5. Monitor rule instances.

See ["About the Monitor Instance-Name Page" on page 586](#).

SEE ALSO

[Add a Rule Instance | 594](#)

[Manage Rule Instances | 596](#)

[Manage Custom Rules | 590](#)

About the Custom KPI Collection Page

IN THIS SECTION

- [Tasks You Can Perform | 578](#)

To access the Custom KPI Collection page, click **Observability > Health > Custom KPI Collection**.

Use this page to view and manage custom rules, and create rule instances. Routing Director uses a custom rule to define key performance indicators (KPIs). You can then apply custom rules to network devices by using rule instances. For more information, see ["Rules Overview" on page 576](#).

Tasks You Can Perform

You can perform the following tasks from this page:

- View rule instances—Click the Rule Instantiation tab to create, modify, and view rule instances. You use a rule instance to apply a custom rule on one or more devices. See ["About the Rule Instantiations Tab" on page 578](#).
- View custom rules—Click the **Custom Rules > Rules List** tab to view a list of all custom rules that you uploaded. See ["About the Rules List Tab" on page 580](#).
- View rule helper files—Click the **Custom Rules > Rule Helper Files** tab to view a list of rule helper files that you uploaded. See ["About the Rule Helper Files Tab" on page 580](#).
- View alerts related to custom rules—Click the Alerts tab to view all alerts associated with the custom rules that you defined. See ["About the Alerts Tab \(Custom Rules\)" on page 584](#).

About the Rule Instantiations Tab

IN THIS SECTION

- [Tasks You Can Perform | 579](#)
- [Field Descriptions | 579](#)

To access this tab, click **Observability > Health > Custom KPI Collection > Rule Instantiations**.

Use this tab to view all rule instances that you created. You create a rule instance to apply a custom rule on one or more devices. For more information on custom rules, see ["Rules Overview" on page 576](#).

Tasks You Can Perform

You can perform the following tasks from this tab:

- View all rule instances.
- Create a rule instance. See ["Add a Rule Instance" on page 594](#).
- Create a new tag or add a tag to a rule instance. See ["Add Tags to a Rule Instance" on page 597](#).
- Edit an existing rule instance. See ["Edit a Rule Instance" on page 596](#).
- Delete an existing rule instance. See ["Delete a Rule Instance" on page 598](#).
- Click a rule instance (hyperlink) to view detail of alerts generated on the Monitor *Rule-Instance* page. See ["About the Monitor Instance-Name Page" on page 586](#).
- Sort, resize, or re-arrange columns in a table (grid).
- Show or hide columns in the table or reset page preferences, using the vertical ellipsis menu.
- Search by using keywords—Click the search icon (magnifying glass), enter the search term in the text box, and press Enter. The search results are displayed on the same page.
- Filter the data displayed in the table—Click the filter icon (funnel) and select whether you want to show or hide advanced filters. You can then add or remove filter criteria, save criteria as a filter, apply or clear filters, and so on. The filtered results are displayed on the same page.

Field Descriptions

[Table 137 on page 579](#) lists the fields on the Rule Instantiations Tab.

Table 137: Fields on the Rule Instantiations Tab

Field	Description
Name	Name of the rule instance.

Table 137: Fields on the Rule Instantiations Tab (*Continued*)

Field	Description
Alert	Alerts generated on the devices added to the rule instance.
Rules	Names of the custom rules added to the rule instance. Hover over + <i>number</i> to view the names of the other rules added to the instance but are not displayed by default.
State	State of the rule instance. The states are: • Running • Paused
Tags	Names of all the tags associated with the rule instance. A tag is a label (unique key/value pair) that you assign to a rule instance. You can use these labels to filter rule instances. Hover over + <i>number</i> to view the names of the other tags added but are not displayed by default.
Description	Brief description of the rule instance, if any.

About the Custom Rules Tab

SUMMARY

Read these topics to understand more about custom rules and rule helper files.

IN THIS SECTION

- [About the Rules List Tab | 581](#)
- [About the Rule Helper Files Tab | 582](#)

To access this tab, click **Observability > Health > Custom KPI Collection > Custom Rules**.

You can access the following tabs from the Custom Rules tab:

- **Rules List Tab**—Use this tab to view all custom rules that you uploaded. See ["About the Rules List Tab" on page 581](#).
- **Rule Helper Files Tab**—Use this tab to view all rule helper files that you uploaded. See ["About the Rule Helper Files Tab" on page 582](#).

About the Rules List Tab

IN THIS SECTION

- [Tasks You Can Perform | 581](#)
- [Field Descriptions | 582](#)

To access this tab, click **Observability > Health > Custom KPI Collection > Custom Rules > Rules List**.

Use this tab to view a list of all custom rules that you uploaded. For more information on custom rules, see ["Rules Overview" on page 576](#).

Tasks You Can Perform

You can perform the following tasks from this tab:

- View all custom rules.
- Upload a custom rule. See ["Upload a Custom Rule" on page 590](#).
- Download a custom rule. See ["Download a Custom Rule" on page 590](#).
- Create a new tag or add a tag to a rule. See ["Add Tags a Custom Rule" on page 590](#).
- Delete a custom rule. See ["Delete a Custom Rule" on page 592](#).
- Sort, resize, or re-arrange columns in a table (grid).
- Show or hide columns in the table or reset page preferences, using the vertical ellipsis menu.
- Search by using keywords—Click the search icon (magnifying glass), enter the search term in the text box, and press Enter. The search results are displayed on the same page.

- Filter the data displayed in the table—Click the filter icon (funnel) and select whether you want to show or hide advanced filters. You can then add or remove filter criteria, save criteria as a filter, apply or clear filters, and so on. The filtered results are displayed on the same page.

Field Descriptions

Table 138 on page 582 lists the fields on the Rules List Tab.

Table 138: Fields on the Rules List Tab

Field	Description
Name	Name of the custom rule.
Instance	Total number of rule instances linked to the custom rule.
Tag	Names of all the tags associated with the rule. A tag is a label (unique key/value pair) that you assign to a rule. You can use these labels to filter rules. Hover over + <i>number</i> to view the names of the other tags added but are not displayed by default.
Description	Brief description of the custom rule, if any.

About the Rule Helper Files Tab

IN THIS SECTION

- [Tasks You Can Perform | 583](#)
- [Field Descriptions | 583](#)

To access this tab, click **Observability > Health > Custom KPI Collection > Custom Rules > Rule Helper Files**.

Use this tab to view a list all rule helper files that you uploaded. You can define how data is fetched from a device in the rule helper file.



NOTE: We recommend that you contact [Juniper Networks Professional Services](#) to define a rule helper file based on your requirements.

For more information see, ["Rules Overview" on page 576](#).

Tasks You Can Perform

You can perform the following tasks from this tab:

- View all rule helper files.
- Upload a rule helper file. See ["Upload a Rule Helper File" on page 593](#).
- Download a rule helper file. See ["Download a Rule Helper File" on page 593](#).
- Delete a rule helper file. See ["Delete a Rule Helper File" on page 593](#).
- Sort, resize, or re-arrange columns in a table (grid).
- Show or hide columns in the table or reset page preferences, using the vertical ellipsis menu.
- Search by using keywords—Click the search icon (magnifying glass), enter the search term in the text box, and press Enter. The search results are displayed on the same page.
- Filter the data displayed in the table—Click the filter icon (funnel) and select whether you want to show or hide advanced filters. You can then add or remove filter criteria, save criteria as a filter, apply or clear filters, and so on. The filtered results are displayed on the same page.

Field Descriptions

[Table 139 on page 583](#) lists the fields on the Rule Helper Files Tab.

Table 139: Fields on the Rule Helper Files Tab

Field	Description
Name	Name of the rule helper file.
Rules	Rules included in the rule helper file. Hover over + <i>number</i> to view the names of the other rules added but are not displayed by default.

About the Alerts Tab (Custom Rules)

IN THIS SECTION

- [Tasks You Can Perform | 584](#)
- [Field Descriptions | 585](#)

To access this tab, click **Observability > Health > Custom KPI Collection > Alerts**.

Use this tab to view alerts generated for devices. These alerts are triggered by the custom rules that you defined. For more information on defining a custom rule, see ["Manage Custom Rules" on page 590](#).

Tasks You Can Perform

You can perform the following tasks from this tab:

- In the widgets, you can view:
 - **Total Alerts**—Displays the total number of alerts generated. The total number includes alerts that are currently open and are addressed (cleared).
 - **Critical Alerts**—Displays the number of active critical alerts that need immediate attention.
 - **Minor Alerts**—Displays the number of active minor alerts that are generated. They are warnings that needs to be fixed but don't require immediate attention.
- View a graphical representation of all active alerts.

On the graph, you can:

- View alerts generated for the past 30 minutes, by default.
- View the total number of active alerts on the y-axis, and the time the alerts occurred on the x-axis.
- View alerts for a week, a day, 3 hours, 1 hour, 30 minutes (default), or a custom time period.
- In the Alerts table, you can:
 - View a device-wise list of all alerts generated for various custom rules.

- View details of an alert—Select an alert and click **More > Detail** or click the **Details** icon on the left of the device name to view more information about the alert. The Alert Details pane appears displaying the details of the alert.
- Sort, resize, or re-arrange columns in a table (grid).

Field Descriptions

Table 140 on page 585 describes the fields in the Alerts table.

Table 140: Fields in the Alert Table

Field	Description
Alert ID	The UUID of the generated alert.
Device	The name of the device.
Rule Instance ID	The UUID of the rule instance.
Severity	The severity level of the alert. The levels are: • Critical—Indicates that the issue needs immediate attention. • Minor—Indicates that the issue is being monitored and currently there is no impact on the functioning of the network or network devices. They are warnings that needs to be fixed but don't require immediate attention.
Raised Time	The date and time that the alert was generated.
Topic	The name of the string used to logically group related rules.
Rule	The custom rule for which the alert is triggered.
Type	The KPIs associated with the custom rule that triggered the alert.
Site	The siite in which the device (for which the alert was raised) is located.

Table 140: Fields in the Alert Table *(Continued)*

Field	Description
Status	The current status of the alert. The statuses are: • Open—An active alert that has not be addressed. • Cleared—The alert is addressed and is no longer active.

About the Monitor Instance-Name Page

SUMMARY

This topic provides information about how you can monitor alerts and key performance indicators (KPIs) that trigger alerts from the **About the Monitor *Instance-Name*** page.

IN THIS SECTION

- [Alerts Tab | 586](#)
- [KPI Tab | 588](#)

To access this page, click **Observability > Health > Custom KPI Collection > Rule Instantiations > *Instance-Name***. Users with the Super User, Network Admin, or Observer roles can use this page.

Routing Director enables you to upload custom rules that define how KPIs for various device components are collected. You can apply a custom rule on devices by creating a rule instance. After you deploy a rule instance, alerts are generated when a KPI crosses the threshold limit defined in the custom rules file.

- You can view these alerts on the **Alerts** tab of the About the Monitor *Rule-Instance* page. These alerts highlight issues that might need attention and how these issues can affect the network.
- You can view more information about changes in KPIs that trigger these alerts on the **KPI** tab of the About the Monitor *Rule-Instance* page.

Alerts Tab

To access this tab, click **Observability > Health > Custom KPI Collection > Rule Instantiations > *Instance-Name* > Monitor *Instance-Name***. The Alerts tab is displayed by default.

Routing Director generates alerts to notify you of anomalies in the KPIs in your network. This tab displays all alerts generated for a rule instance.

Tasks You Can Perform

You can perform the following tasks from this tab:

- In the widgets, you can view:
 - Total Alerts—Displays the total number of alerts generated for all the KPIs associated with a device in the rule instance. The total number includes alerts that are currently open and are addressed (cleared).
 - Critical Alerts—Displays the number of active critical alerts that need immediate attention.
 - Minor Alerts—Displays the number of active minor alerts that are generated. They are warnings that needs to be fixed but don't require immediate attention.
- View a graphical representation of all active alerts.

On the graph, you can:

- View alerts generated for the past 30 minutes, by default.
- View the total number of active alerts on the y-axis, and the time the alerts occurred on the x-axis.
- View alerts for a week, a day, 3 hours, 1 hour, 30 minutes (default), or a custom time period.
- In the Alerts table, you can:
 - View a device-wise list of all alerts generated for various KPIs associated with the rule instance.
 - View details of an alert—Select an alert and click **More > Detail** or click the **Details** icon on the left of the device name to view more information about the alert. The Alert Details pane appears displaying the details of the alert.
 - Sort, resize, or re-arrange columns in a table (grid).

Fields Descriptions

[Table 141 on page 587](#) describes the fields in the Alerts table.

Table 141: Fields in the Alerts Table

Field	Description
Alert ID	The UUID of the generated alert.

Table 141: Fields in the Alerts Table (*Continued*)

Field	Description
Device	Name of the device.
Rule Instance ID	The UUID of the rule instance.
Severity	Severity level of the alert. The levels are: • Critical—Indicates that the issue needs immediate attention. • Minor—Indicates that the issue is being monitored and currently there is no impact on the functioning of the network or network devices. They are warnings that needs to be fixed but don't require immediate attention.
Raised Time	Date and time that the alert was generated.
Topic	The name of the string used to logically group related rules.
Rule	The custom rule for which the alert is triggered.
Type	The KPIs associated with the custom rule that triggered the alert.
Site	Site in which the device (for which the alert was raised) is located.
Status	The current status of the alert. The statuses are: • Open—An active alert that has not be addressed. • Cleared—The alert is addressed and is no longer active.

KPI Tab

To access this tab, click **Observability > Health > Custom KPI Collection > Rule Instantiations > *Instance-Name* > Monitor *Instance-Name* > KPI (tab)**.

Routing Director generates alerts based on the anomalies in the KPIs in your network. You can view a graphical representation of the performance, and alerts generated for all KPIs associated with a device on this page. You can set triggers and threshold limits for individual KPIs when you define rule

parameters. You configure these rule parameters when you define custom rules. To define a custom rule based on your requirements, we recommend that you contact [Juniper Networks Professional Services](#).

Tasks You Can Perform

You can perform the following tasks from this page:

- View graphical representation of the performance, and alerts generated for all KPIs associated with a device. To do this information:
 1. Select a device from the **Device Selection** drop-down list.
 2. Select a rule from the **Rule Selection** drop-down list.

The rule parameters (triggers, threshold limits) associated with the selected rule and device are displayed in accordions on the tab. Click each accordion to view a graph of the performance, and alerts associated with each rule parameter. For more information, see .



NOTE: The name of the accordion is the name of the rule parameters that you set when you defined the custom rule.

- You can view KPI information for a week, a day, 3 hours, 1 hour, 30 minutes (default), or a custom time period.
- View the time stamp of the next refresh (in DD/MM/YYYY hh:mm:ss AM/PM format) in the top-right corner of the tab. For Example: 07/12/2024 12:42:43 PM. The graph auto-refreshes at an interval depending on the time range for which the information is displayed. See [Table 142 on page 589](#) for more information.
- Update the details on the tab by clicking the refresh icon.

Table 142: Auto-Refresh Rate

Time Range	Auto-Refreshed
Weekly	Every 16.8 hours
Daily	Every 58 minutes
Every three hours	Every 8 minutes
Hourly	Every 3 minutes

Table 142: Auto-Refresh Rate *(Continued)*

Time Range	Auto-Refreshed
Thirty minutes	Every 2 minutes
Custom	No auto-refresh

Manage Custom Rules

SUMMARY

Read these topics to understand how to manage custom rules.

IN THIS SECTION

- [Upload a Custom Rule | 590](#)
- [Download a Custom Rule | 591](#)
- [Add Tags to a Custom Rule | 591](#)
- [Delete a Custom Rule | 592](#)

You can upload a custom rule in JSON format, download a rule, add tags to a rule, and delete a rule from the **Observability > Health > Custom KPI Collection > Custom Rules > Rule List** tab.



NOTE: Before you can upload a custom rule by using the Routing Director GUI, you must define a custom rule. We recommend that you contact [Juniper Networks Professional Services](#) to define a custom rule based on your requirements.

Upload a Custom Rule

You use a custom rule to define key performance indicators (KPIs). You can organize these rules into rule instances and apply the rules to devices. You can upload a custom rule in the JSON format from the Rule List tab.

To upload a custom rule:

1. Navigate to the Rules List tab.
2. Click **Upload Rule**.

The Upload Rule page appears.



NOTE: Fields marked with an asterisk (*) are mandatory.

3. Click **Browse** and navigate to location of the rules file on your local file system.

4. Click **Open** after you select the rules file.

The Upload Rule page appears.

5. Click **Upload** to upload the rule.

The Rule List tab appears and displays the newly added custom rule.

Download a Custom Rule

You can download an existing custom rule in the JSON format from the Rule List Tab.

To download a custom rule:

1. Navigate to the Rules List tab.
2. Select the custom rule that you want to download by clicking the check box next to the name of the rule, and click **Download**.

The rule is downloaded and saved in your local file system.

Add Tags to a Custom Rule

You can add tags to a custom rule. A tag is a label (unique key/value pair) that you assign to a rule. You can use these labels to filter rules.

To add a tag to a custom rule:

1. Navigate to the Rules List tab.
2. Select the rule that you want to add a tag to by selecting the check box next to the name of the rule.
3. Click **More > Manage Tags**.

The Manage Tag page appears.

4. Select the tags that you want to add to the rule from the **Tags** drop-down list.

Alternatively, you can also create a new tag. To create a new tag:

- a. Click **Create New Tag +**.

The Create Tag page appears.

- b. Enter a name for the tag in the **Tag** field.

You must add tags in the key:value format where both are text separated by a colon. A key must start with a letter and must be followed by letters, numbers, periods (.), hyphens (-), or underscores (_). A key must be between 1 to 200 characters long. A value can include alphanumeric characters, periods (.), hyphens (-), and underscores (_). A value must be between 1 to 200 characters long.

- c. Click **OK** to create the new tag.

The new tag is automatically added to the **Tags** drop-down list on the Manage Tags page.

5. Click **OK** on the Manage Tags page to add the tag(s) to the rule.

The Rules List tab appears and displays the newly added tags in the rule.

Delete a Custom Rule

You can delete a custom rule from the Rules List tab. Before you delete a rule, you must ensure that the custom rule that you want to delete is not added to any rule instance.

To delete a rule instance:

1. Navigate to the Rules List tab.
2. Select the custom rule that you want to delete by selecting the check box next to the name of the rule, and click the **Delete** (Dustbin) icon.

The Confirm Delete pop-up appears.

3. Click **Yes** to delete the rule.

The rule is deleted and removed from the All Custom Rules table.

SEE ALSO

[Rules Overview](#) | 576

Manage Rule Helper Files

SUMMARY

Read these topics to understand how to manage rule helper files.

IN THIS SECTION

- [Upload a Rule Helper File | 593](#)
- [Download a Rule Helper File | 594](#)
- [Delete a Rule Helper File | 594](#)

You use a rule helper file to define how data is fetched from a device. You can upload, download, and delete a rule helper file from the **Observability > Health > Custom KPI Collection > Custom Rules > Rule Helper Files** tab.

Upload a Rule Helper File

You can upload a rule helper file to define how data is fetched from a device.



NOTE: We recommend that you contact [Juniper Networks Professional Services](#) to define a rule helper file based on your requirements.

To upload a rule helper file:

1. Navigate to the Rule Helper Files tab.
2. Click **Upload Rule**.

The Upload Helper File page appears.



NOTE: Fields marked with an asterisk (*) are mandatory.

3. Click **Browse** and navigate to location of the file on your local file system.
4. Click **Open** after you select the file.

The Upload Helper File page appears.

5. Click **Upload** to upload the rule helper file.

The Rule Helper Files tab appears and displays the newly added rule helper file.

Download a Rule Helper File

You can download an existing rule helper file from the Rule List Tab.

To download a rule helper file:

1. Navigate to the Rule Helper Files tab.
2. Select the file that you want to download by clicking the check box next to the name of the file, and click **Download**.

The rule helper file is downloaded and saved in your local file system.

Delete a Rule Helper File

You can delete a rule helper file from the Rules Helper files tab. Before you delete a rule helper file, you must ensure that the file that you want to delete is not associated with any rule.

To delete a rule instance:

1. Navigate to the Rules Helper Files tab.
2. Select the rule helper file that you want to delete by selecting the check box next to the name of the file, and click the **Delete** (Dustbin) icon.

The Confirm Delete pop-up appears.

3. Click **Yes** to delete the rule helper file.

The rule helper file is deleted and removed from the All Helper Files table.

SEE ALSO

| [Rules Overview](#) | [576](#)

Add a Rule Instance

Routing Director uses a custom rule to define key performance indicators (KPIs). You can then apply custom rules to network devices by using rule instances. For more information, see ["Rules Overview" on page 576](#).

You can create a rule instance by using the Routing Director UI. However, before you create a rule instance, you must upload a custom rule. For more information, see ["Upload a Custom Rule" on page 590](#).

To add a rule instance:

1. Click **Observability > Health > Custom KPI Collection > Rule Instantiations**.

The All Rule Instantiations tab appears.

2. Click the Add (+) icon.

The General section of the Create Instance wizard appears.



NOTE: Fields marked with an asterisk (*) are mandatory.

3. Enter the following information in the General section:

- a. Enter a name for the instance in the **Name** field.
- b. Enter a description for the instance in the **Description** field.
- c. Select rules that you want to add to the instance. To select a rule:
 - i. Click the add (+) icon in the Selected Rules section. The Select Rules page appears.
 - ii. Select all the rules that you want to add to the instance in the Selected Rules section by clicking the check boxes next to the name of the rules.
 - iii. Click **OK** to confirm selection.

4. Click **Next**.

The Devices section of the Create Instance wizard appears.

5. Select all devices that you want to add to the instance in the Device section.

To select a device:

- a. Click the add (+) icon in the Selected Devices section. The Select Devices page appears.
- b. Select all the devices that you want to add to the instance by clicking the check boxes next to the name of the devices.
- c. Click **OK** to confirm selection.

6. Click **Next**.

The Parameters section of the Create Instance wizard appears. You can view all the rules that you selected in step ["3.c" on page 595](#) in this section.

7. Click **Next** to view a summary of the general information, rules added, devices selected, and rule parameters that you provided in the various sections of the wizard.

8. Do any of the following:

- a. Click **Cancel** to exit the wizard without saving changes.
- b. Click **Back** to go to the Parameters section of the Create Instance wizard.
- c. Click **Execute** to create the rule instance.

The Monitor Execute Status page appears displaying the overall progress and the final status (success or failure) of the rule instance execution job.

9. After the rule instance execution job is completed, click **Close & Exit**.

The Rules Instantiations tab appears and displays the newly created rule instance.

Manage Rule Instances

SUMMARY

Read these topics to understand how to manage rule instances.

IN THIS SECTION

- [Edit a Rule Instance | 596](#)
- [Add Tags to a Rule Instance | 597](#)
- [Delete a Rule Instance | 598](#)

You can edit a rule instance, add tags to a rule instance, and delete a rule instance from the Rule Instantiations tab.

Edit a Rule Instance

To edit an existing rule instance:

1. Click **Observability > Health > Custom KPI Collection > Rule Instantiations**.

The All Rule Instantiations tab appears.

2. Select the rule that you want to edit by selecting the check box next to the name of the rule instance, and click the **Edit** (Pen) icon.

The Modify *Instance Name* page appears displaying the Create Instance wizard.

3. Edit the information in the General, Devices, and Parameters section of the Create Instance wizard as required.

4. Do any of the following:

- a. Click **Cancel** to exit the wizard without saving changes.
- b. Click **Back** to go to the Parameters section of the Create Instance wizard.
- c. Click **Execute** to create the rule instance.

The Monitor Execute Status page appears displaying the overall progress and the final status (success or failure) of the rule instance execution job.



NOTE: The **Execute** button is disabled if there are no changes made to the rule instance.

5. After the rule instance execution edit job is completed, click **Close & Exit**.

The Rules Instantiations tab appears and displays the edited rule instance.

Add Tags to a Rule Instance

You can add tags to an existing rule instance. A tag is a label (unique key/value pair) that you assign to a rule instance. You can use these labels to filter rule instances.

To add a tag to a rule instance:

1. Click **Observability > Health > Custom KPI Collection > Rule Instantiations**.

The All Rule Instantiations tab appears.

2. Select the rule instance that you want to add a tag to by selecting the check box next to the name of the rule instance.

3. Click **More > Manage Tags**.

The Manage Tag page appears.

4. Select the tags that you want to add to the instance from the **Tags** drop-down list.

Alternatively, you can also create a new tag. To create a new tag:

a. Click **Create New Tag +**.

The Create Tag page appears.

b. Enter a name for the tag in the **Tag** field.

You must add tags in the key:value format where both are text separated by a colon. A key must start with a letter and must be followed by letters, numbers, periods (.), hyphens (-), or underscores (_). A key must be between 1 to 200 characters long. A value can include

alphanumeric characters, periods (.), hyphens (-), and underscores (_). A value must be between 1 to 200 characters long.

- c. Click **OK** to create the new tag.

The new tag is automatically added to the **Tags** drop-down list on the Manage Tags page.

5. Click **OK** on the Manage Tags page to add the tag(s) to the rule instance.

The Rules Instantiations tab appears and displays the newly added tags in the rule instance.

Delete a Rule Instance

You can delete a rule instance from the Rule Instantiations tab. To delete a rule instance:

1. Click **Observability > Health > Custom KPI Collection > Rule Instantiations**.

The All Rule Instantiations tab appears.

2. Select the rule instance that you want to delete by selecting the check box next to the name of the rule instance, and click the **Delete** (Dustbin) icon.

The Confirm Delete pop-up appears.

3. Click **Yes** to delete the rule instance.

The rule instance is deleted and removed from the All Rule Instantiations table.

View KPI Performance and Alerts (Custom Rules)

SUMMARY

Routing Director enables you to upload custom rules that define how key performance indicators (KPIs) for various device components are collected. This topic provides information on how you can view and monitor KPI performance, and alerts that are triggered when a KPI crosses the threshold limit that you defined.

IN THIS SECTION

- [Overview | 598](#)
- [Components on the KPI Tab | 599](#)
- [KPI Performance Graphs | 599](#)

Overview

Routing Director generates alerts based on the anomalies in the KPIs in your network. You can view a graphical representation of the performance, and alerts generated for all KPIs associated with a device

on the KPI tab (**Observability > Health > Custom KPI Collection > Rule Instantiations > *Instance-Name* > Monitor *Instance-Name* > KPI (tab)**)).

You can set triggers and threshold limits for individual KPIs when you define rule parameters. You configure these rule parameters when you define custom rules. To define a custom rule based on your requirements, we recommend that you contact [Juniper Networks Professional Services](#).

Components on the KPI Tab

You can view the following components on the KPI tab.

- **Device Selection and Rule Selection drop-down lists**—To view KPI performance and alerts, you must first select a device from the **Device Selection** drop-down list. After you select a device, you must select a rule (applied on the selected device) from the **Rule Selection** drop-down list. The rule parameters (or KPIs) associated with the selected rule and device are displayed as accordions on the tab.
- **Accordions on the page**—The accordions on the page are displayed based on the device and rule that you selected. The name of the accordion is the name of the rule parameters that you set when you defined the custom rule. For more information on the accordions displayed on the page, see "[KPI Performance Graphs](#)" on page 599.
- **Time Period bar**—You can view KPI information for a week, a day, 3 hours, 1 hour, 30 minutes (default), or a custom time period. Use the Time Period bar in the top-right corner of the tab to change the time period settings.
- **Refresh icon**—Click the refresh icon to update the details on the tab immediately. If you do not use the refresh icon, the page auto-refreshes at an interval depending on the time range for which the information is displayed. The next auto refresh is displayed in DD/MM/YYYY hh:mm:ss AM/PM format just below the Time Period bar.

KPI Performance Graphs

The graphs on the KPI Tab display the performance and alerts associated with a KPI. The rule parameters (or KPIs) associated with the selected rule and device are displayed as accordions on this tab. The name of the accordion is the name of the rule parameters that you set when you defined the custom rule.

Each accordion on this tab has three sections.

- The first section displays the name of the accordion, and details of the device and rules.
- The second section displays the alerts generated. Critical alerts are displayed in red and minor in yellow.
- The last section displays a performance graph representing the actual performance of the KPI (associated with the selected device) for the past 30 minutes.

Click any accordion to expand it. The highest severity alert associated with the KPI is displayed under the name of the accordion. The device's keys are listed below the alerts in the order of severity of the events that have occurred on them. You can view up to six keys at a time with the key that is in the most critical state displayed at the top of the list. To view keys that are not listed, click the **Select Keys** drop-down list and select the interface. However, you must clear a previously selected key to be able to select another key.

Click the toggle button next to the name of an key, to view the performance of that interface in a graph. Depending on the KPI, the graph can either be a line graph or a state graph. You can also zoom into a particular portion of the graph to view more information about events that have occurred. To view details of more than one interface on the graph, click the toggle button next to the name of the interface in the Show Keys list. Details of the interfaces displayed on the graph are also listed in a table below the graph. You can also click the option buttons on the left of an interface's name in the table to highlight the performance for that interface on the graph.

You can view performance and alerts for a week, a day, 3 hours, 1 hour, 30 minutes, or a custom time period. By default, performance for the past 30 minutes is displayed. To change this period, click the **Week, Day, 3 hrs, 1 hr, 30 mins**, or **Custom** buttons provided above the graph.

You can view more than 25 data points on a graph related to events (in real time) that have occurred on the interface when you select the 30-minute time period. However, you can only view up to 25 data points related to events when you select a week, a day, 3 hours, 1 hour, or a custom time period (of more than 30 minutes). Data is aggregated to ensure that not more than 25 data points are plotted on the graph at once.

The graph auto-refreshes at an interval depending on the time range for which the information is displayed. See [Table 143 on page 601](#) for more information. However, you can refresh the graph at any point by clicking the **Refresh** icon provided above the graph. You can also click the pop-out button next to the graph to open the graph in a new tab. You can view all customizations that you made on the graph in the parent tab, in the new tab.

Alerts, if any, related to the interface is also displayed on the graph, and in the table below the graph. You can also open the graph in a new tab. When you open a graph in a new tab, you can view the following information in the new tab as well:

- Alerts related to the interface.

Alerts are refreshed across all open tabs simultaneously, when:

- An alert is refreshed in any one of the open tabs.
- The last alert fetched was beyond three minutes.
- List of interfaces that you toggled to view from the Select Keys drop-down list.
- Interface that you selected from the table below the graph.

Table 143: Auto-Refresh Rate

Time Range	Auto-Refreshed
Weekly	Every 16.8 hours
Daily	Every 58 minutes
Every three hours	Every 8 minutes
Hourly	Every 3 minutes
Thirty minutes	Every 2 minutes
Custom	No auto-refresh

You can similarly view the graphs and alerts related to the performance of other KPIs associated with the device and rule that you selected.

5

PART

Trust and Compliance

- Introduction | **603**
 - Manage Trust Settings and Trust Scores | **607**
 - Manage Compliance Scans | **623**
 - Manage Vulnerabilities | **631**
 - Monitor Integrity | **638**
-

Introduction

IN THIS CHAPTER

- [Trust and Compliance Overview | 603](#)
- [Perform Compliance Scan and Manage Checklists | 604](#)

Trust and Compliance Overview

IN THIS SECTION

- [Benefits | 604](#)

As enterprises and service providers scale up their network infrastructure to meet the increasing connectivity needs of subscribers, their networks become increasingly complex because of the number of devices that connect to the network. Service providers must meet the connectivity and bandwidth requirements of mobile, IoT, and other devices that connect daily, while keeping the network secure. Possibilities of threats that can lead to a network outage from devices that connect to the network highlight the need to proactively address device and network security concerns. Service providers need to ensure that connectivity is uninterrupted without impacting security.

Routing Director helps protect the devices and the network as a whole by taking the principle of zero trust networking (ZTN) to the next level. Zero trust security considers all devices, whether within or outside the network, as untrusted. Routing Director extends this concept by periodically evaluating the device's configuration, integrity, and performance against standards applied on the network and recommends corrective measures to keep the network secure.

Routing Director assigns a trust score to each target. The term target refers to a device or a device component. This trust score is generated based on the integrity of the software and hardware components, vulnerabilities defined in SIRT advisories, and compliance with rules defined in the benchmarks document applied to the network. A benchmarks document contains recommendations and

baseline configurations for securely configuring software, devices, and network infrastructure. Depending on changes in the network, Routing Director continually updates the trust score.

In addition to the trust score assigned to devices, Routing Director also alerts you when a device doesn't comply with the rules in the benchmarks document; for example, when a device or the OS running on a device reaches its EOL. You can view this information on the Alerts tab (**Observability > Health > Events > Alerts**) as well as on the **Observability > Health > Troubleshoot Devices > *Device-Name* > Overview** tab.

Routing Director also provides a dashboard (**Observability > Health > Health Dashboard > Trust**) that enables you to view the overall trust score. You can also monitor in real-time, vulnerabilities that affect targets in the network, view device integrity, and the compliance status of devices. You can view a graph of the average trust score, the trust plan applied, and details of unhealthy devices through a single pane of glass GUI.

Benefits

Routing Director protects the network by:

- Continuously monitoring the targets and providing information about potential vulnerabilities.
- Measuring trustworthiness of the devices on the network by assigning a trust score to each network target.
- Providing information to perform corrective action on non-compliant devices.
- Providing a dashboard to view and monitor, in real-time, device integrity, compliance, and vulnerabilities that affect targets enabling you to perform corrective action.

RELATED DOCUMENTATION

[Integrity of the Hardware and Software on the Network | 638](#)

[Vulnerabilities Overview | 631](#)

[Compliance Standards Overview | 607](#)

[Trust Score Overview | 621](#)

Perform Compliance Scan and Manage Checklists

Routing Director enables network administrators to measure the trustworthiness of your network. It performs periodical scans and alerts you when a target does not meet the requirements that are specified in the applied benchmarks document.

Scanning and monitoring individual targets in a large network for compliance with industry-accepted standards or benchmarks can be challenging and time-consuming. Routing Director addresses this challenge automating the process of checking the compliance of your networks and targets using compliance checklists. Automating the process saves time and reduces the risk of errors that can result from manual checking. Routing Director allows you to update compliance checklists by importing rule results from completed scans. These rule results are generated from a benchmarks document and indicate a target's compliance with the rules defined in the benchmarks document. An administrator can edit the status of the rules in the checklist depending on the requirements of the network.

By default, Routing Director runs automated compliance scans every 24 hours. However, you can run a custom compliance scan at any time to assess the trust posture of targets. [Figure 45 on page 605](#) shows the sequence of tasks that you must perform to run compliance scans and manage compliance checklists.

Figure 45: Run compliance scans and manage compliance checklists



The workflow for running a compliance scan and updating compliance checklists is as follows:

1. Run a compliance scan by selecting a benchmarks document, tailorings document, and the targets that you want to scan. You can also run a scan by uploading a preconfigured compliance configuration file in the JSON format.

You can run a scan for a single device or for multiple devices at a time. After the scan is completed, Routing Director generates a report of compliance of targets with the applied benchmarks document. See ["Perform Custom Compliance Scans" on page 626](#).

2. From the Compliance Checklist (**Settings > Compliance Checklist**) page, create a compliance checklist by specifying a checklist name and the target to which you want to apply the checklist. See ["Add Checklist for a Device" on page 616](#).

The checklist appears on the Compliance Checklist page.

3. Import scan results and update the checklist:

- a. Import scan results to the checklist that you created.
- b. Review the statuses of rules in the Status column and update the status as needed.
- c. Save the checklist.

See ["Import Scans and Update Rule Results in a Checklist" on page 617](#)

Routing Director automatically uses this checklist when scanning the target in the future.

Manage Trust Settings and Trust Scores

IN THIS CHAPTER

- Compliance Standards Overview | 607
- About the Compliance Benchmarks Page | 608
- About the Compliance Tailorings Page | 610
- Example: Create a Tailoring Document for NTP Settings | 612
- About the Compliance Checklist Page | 613
- Add a Checklist Template | 615
- Add Checklist for a Device | 616
- Import Scans and Update Rule Results in a Checklist | 617
- Trust Plans Overview | 617
- About the Network Score Formula Page | 619
- Trust Score Overview | 621

Compliance Standards Overview

Routing Director follows the compliance standards and specifications defined by the National Institute of Standards and Technology (NIST), specifically the Security Content Automation Protocol (SCAP). Compliance documents follow the Extensible Configuration Checklist Description Format (XCCDF) specification defined using SCAP by NIST.

SCAP (*pronounced ess-cap*) is a suite of specifications for exchanging security automation content used to assess configuration compliance and to detect vulnerable versions of software. Multiple tools can use the same SCAP content to perform an assessment that the content describes.

The SCAP languages provide standard vocabularies and conventions for expressing security policy, technical check mechanisms, and assessment results. For more information about SCAP, see [Security Content Automation Protocol](#) at the NIST website.

Of the number of specifications available within the languages category, the XCCDF and the Open Vulnerability and Assessment Language (OVAL) are the primary specifications that are used in Routing Director.

XCCDF is an XML-based specification for writing security checklists, benchmarks, and related documents. An XCCDF document represents a structured collection of security configuration rules for a set of target systems.

The specification is designed to support information interchange, document generation, organizational and situational tailoring, automated compliance testing, and compliance scoring.

Routing Director monitors the devices and software to ensure that they comply with the security rules that are defined in the benchmarks and tailorings documents applied to the network.

RELATED DOCUMENTATION

[Compliance Scans Overview](#) | 623

[About the Compliance Benchmarks Page](#) | 608

About the Compliance Benchmarks Page

IN THIS SECTION

● [Tasks You Can Perform](#) | 609

To access this page, click **Trust > Compliance > Benchmarks**.

Routing Director automatically monitors the hardware and software in the network for compliance with the rules defined in the benchmarks documents. Benchmarks documents consist of compliance policies and rules defined by the Center for Internet Security (CIS). CIS benchmarks help you protect devices, software, and networks from cyber threats. The benchmarks document contains profiles which are based on the policies that are defined in the document. Profiles determine how rules and policies are enforced on the network devices to obtain the desired level of compliance. You can apply one or more profiles to your network devices. To view the latest benchmark documents, visit the [CIS Benchmarks](#) page.

The Compliance Benchmarks page lists the rules that are defined in the selected benchmarks document. Along with the rules, you can also view information about the actions that you can take if a device does

not comply with the rules that are defined in the benchmarks document. Click the Details icon for the rule to view more information about the rule and the action to be taken if a target doesn't comply with the rule. A typical benchmark document contains two predefined profiles – Level 1 and Level 2. The Level 1 profile is the base recommendation that doesn't cause much performance impact. The Level 2 profile is meant for environments, such as defense systems, where security is of utmost importance and can sometimes impact performance if due care is not taken during implementation. If no profile is selected the profile <default> is selected.



NOTE: We recommend that you perform a test implementation before implementing Profile 2 in production environments.

Tasks You Can Perform

You can perform the following tasks from this page:

- View details of the benchmarks document applied on the network. Select a profile to view the rules defined for that profile. Also, you can click the **Details** icon to view details about a benchmark document.
- You can also perform the following tasks from this page:
 - Sort, resize, or re-arrange columns in a table (grid).
 - Show or hide columns in the table or reset page preferences, using the vertical ellipsis menu.
 - Search by using keywords—Click the search icon (magnifying glass), enter the search term in the text box, and press Enter. The search results are displayed on the same page.
 - Filter the data displayed in the table—Click the filter icon (funnel) and select whether you want to show or hide advanced filters. You can then add or remove filter criteria, save criteria as a filter, apply or clear filters, and so on. The filtered results are displayed on the same page.
 - Set the number of rows to be displayed in the table—Set the number of rows (15, 30, 50, or 100) to be displayed in the Network Implementation Plan table in the **Display** drop-down list provided at the bottom right-corner of the table.

Use the arrow heads to traverse through the pages or enter a page number to view the entries listed on that page.

About the Compliance Tailorings Page

IN THIS SECTION

- [Tasks You Can Perform | 610](#)
- [Fields on the Compliance Tailorings Page | 611](#)

To access this page, click **Trust > Compliance > Tailorings**.

The XCCDF specification defines tailoring as an element that specifies profiles to modify the behavior of a benchmark. Tailoring is the process of customizing the benchmarks document before you assess the targets in the network. You can create customized tailoring documents for one or more rules defined in the benchmark document. Tailoring documents contain the rules and parameters that the devices on the network should comply with. For example, you can create a tailoring document to define NTP settings for your network. If the network doesn't comply with the parameters that are defined in the tailoring document for NTP settings, Routing Director flags the target as non-compliant with this rule.

The Compliance Tailorings page displays the tailorings documents that are applied on the network.

Tasks You Can Perform

You can perform the following tasks from this page:

- View available tailoring documents. Click the **Details** icon to view more information about the tailoring document.
- Add a new tailoring document. See ["Example: Create a Tailoring Document for NTP Settings" on page 612](#)
- Delete a tailoring document.
- You can also perform the following tasks from this page:
 - Sort, resize, or re-arrange columns in a table (grid).
 - Show or hide columns in the table or reset page preferences, using the vertical ellipsis menu.
 - Search by using keywords—Click the search icon (magnifying glass), enter the search term in the text box, and press Enter. The search results are displayed on the same page.

- Filter the data displayed in the table—Click the filter icon (funnel) and select whether you want to show or hide advanced filters. You can then add or remove filter criteria, save criteria as a filter, apply or clear filters, and so on. The filtered results are displayed on the same page.
- Set the number of rows to be displayed in the table—Set the number of rows (15, 30, 50, or 100) to be displayed in the Network Implementation Plan table in the **Display** drop-down list provided at the bottom right-corner of the table.

Use the arrow heads to traverse through the pages or enter a page number to view the entries listed on that page.

Fields on the Compliance Tailorings Page

Table 144 on page 611 describes the fields on the Compliance Tailorings page.


NOTE: Fields marked with an asterisk (*) are mandatory.

Table 144: Fields on the Compliance Tailorings Page

Field	Description
Name*	Name of the tailorings document.
Version*	Version of the tailorings documents. For example, 1.0.0.
Description	Description of the tailorings document.
Source	Source of the tailorings document. For example: Center for Internet Security (CIS). Routing Director uses the benchmarks document that is defined by CIS.
Benchmark Name	Name of the benchmark as defined by CIS. For example, Juniper OS is the name of the benchmark document that CIS defined for Juniper Networks devices.

Table 144: Fields on the Compliance Tailorings Page (*Continued*)

Field	Description
Benchmark Version	Version of the benchmark document as defined by CIS. For example, v2.1.0a is the version of the benchmark document that CIS defined for Juniper Networks devices.
Profile	The profile selected for the tailoring document.

RELATED DOCUMENTATION

[Example: Create a Tailoring Document for NTP Settings](#) | 612

Example: Create a Tailoring Document for NTP Settings

This topic shows how to create a custom tailorings document for NTP settings.

To create a tailoring document:

1. Click the Add (+) icon on the Compliance Tailorings page.
The Create Tailoring Document wizard appears.
2. On the Select Benchmark page, select a Profile from the list.
As the benchmark document is already applied, the Source, Benchmark, and the Version fields are prepopulated.
3. Click **Next**.
4. (Mandatory) In the Properties page, enter a name and a version number for the tailoring document, and then click **Next**.
In the Tailor Values page, the wizard displays all the default parameters (retrieved from the benchmark document) required for a tailoring document. You can edit the values as necessary. For example, you can edit the NTP server IP address and click the check mark (✓) to save it.
5. Click **Next** to view a summary of the changes.
6. Click **Create**.

A confirmation message, *Tailoring document created successfully.*, is displayed and you can see the new tailoring document displayed on the Compliance Tailorings page.

You can create multiple tailoring documents that are based on the needs of your network. Routing Director uses these tailoring documents to generate compliance reports of devices on the network.

About the Compliance Checklist Page

IN THIS SECTION

- [Tasks You Can Perform | 613](#)
- [Fields on the Compliance Checklist Page | 614](#)

To access this page, click **Trust > Compliance > Checklist**.

A checklist is an organized collection of rules that can be applied to a specific target. Checklists are based on the benchmarks document that is applied to the network. Currently, Routing Director supports only the CIS Juniper OS Benchmarks document.

You use a checklist to assess the compliance of a target against the benchmarks document applied to the network.

Tasks You Can Perform

You can perform the following tasks from this page:

- Add custom checklists for specific targets in the network. See ["Add Checklist for a Device" on page 616](#).
- View available checklists. You can also view and edit the rules that are defined in the checklists.
- View and add checklist templates. See ["Add a Checklist Template" on page 615](#).
- Delete checklists.
- You can also perform the following tasks from this page:
 - Sort, resize, or re-arrange columns in a table (grid).
 - Show or hide columns in the table or reset page preferences, using the vertical ellipsis menu.
 - Search by using keywords—Click the search icon (magnifying glass), enter the search term in the text box, and press Enter. The search results are displayed on the same page.
 - Filter the data displayed in the table—Click the filter icon (funnel) and select whether you want to show or hide advanced filters. You can then add or remove filter criteria, save criteria as a filter, apply or clear filters, and so on. The filtered results are displayed on the same page.

- Set the number of rows to be displayed in the table—Set the number of rows (15, 30, 50, or 100) to be displayed in the Network Implementation Plan table in the **Display** drop-down list provided at the bottom right-corner of the table.

Use the arrow heads to traverse through the pages or enter a page number to view the entries listed on that page.

Fields on the Compliance Checklist Page

Table 145 on page 614 describes the fields on the Compliance Checklist page.

Table 145: Fields on the Compliance Checklist Page

Column	Description
ID	Unique identifier that Routing Director generates for the checklist.
Name	User-defined name for the checklist.
Labels	Labels that were assigned while creating the checklist. You use labels to easily filter and identify the checklist.
Description	Detailed description of the checklist.
Target	Target to which checklist is applied.
Hostname	Hostname of the target.
IP Address	IP address of the target.
Friendly Name	Name used to easily identify the device.
Template	Checklist template used to create the checklist.
Imported Scans	List of scans imported to the checklist.
Open	Displays the compliance status of targets when the rule is under review.

Table 145: Fields on the Compliance Checklist Page *(Continued)*

Column	Description
Not Applicable	Displays the compliance status of targets when the rule does not apply to the target.
Resolved	Displays the compliance status of targets when the rule is resolved.
Total	Displays the number of targets to which checklist is applied.
Last Updated	The time the checklist document was last updated and saved.

RELATED DOCUMENTATION

[Add Checklist for a Device | 616](#)

[Import Scans and Update Rule Results in a Checklist | 617](#)

Add a Checklist Template

Routing Director allows you to create checklist templates, which you can reuse when you create compliance checklists. A checklist template is based on a benchmarks document and you can add multiple checklists that are based on the same checklist template. You use a checklist template to easily create device specific checklists.

To create a checklist template:

1. Click **Trust > Compliance > Checklist** and select the **Templates** tab.
2. On the Templates page, click **Add (+)**.
The Add Checklist Template page appears.
3. (Mandatory) Enter a name and version for the checklist, and select the benchmarks document on which the template should be based.
4. Click **OK**.

The Checklist Template is created and displayed in the **Templates** tab of the Compliance Checklist page.

RELATED DOCUMENTATION

[About the Compliance Checklist Page | 613](#)

[Add Checklist for a Device | 616](#)

Add Checklist for a Device

Checklists are documents that provide guidelines and recommendations for securing networks. A network administrator uses checklists to ensure that the targets and the network meet the security and compliance requirements. Checklists serve as a reference document for network administrators to compare the current configuration of the target to the configuration recommended in the checklist.

A checklist is based on a benchmarks document and contains a set of rules imported from previous scan results. As a checklist might contain hundreds of rules, analyzing and resolving each failed rule for every scan can be a time-consuming task. Routing Director enables you to supplement rule results from scans manually, and allows you to specify that a rule doesn't apply to a specific device.

Routing Director allows you to add a checklist for a specific device, and update it by importing rules from completed scans. You can then edit and mark rules as **Resolved** or **Not Applicable** based on network and device requirements. After you mark a rule as **Resolved** or **Not Applicable**, Routing Director maintains a record of these changes so that a network administrator knows that the rule has been reviewed.

This topic describes how to add a checklist for a device.

1. Click **Trust > Compliance > Checklist**.

A list of existing checklists is displayed.

2. Click **Add (+)**.

3. (Mandatory) On the Add Checklist page, enter a name, select a checklist template, and the device to which the checklist is applied.

4. Click **OK**.

The new checklist is displayed on the Compliance Checklist page.

RELATED DOCUMENTATION

[About the Compliance Checklist Page | 613](#)

Import Scans and Update Rule Results in a Checklist

Routing Director enables you to update checklists by importing rules from a previous scan to an existing checklist and use that checklist for planning manual updates to the device. While you import results from a scan, you can customize them based on the security requirements of your network. If a rule in the scan results is not relevant to the target, you can change the status of the rules to **Not Applicable**. Alternatively, you can manually resolve the rule that failed by setting its status to **Resolved**.

This topic describes how to update a checklist by importing rules from existing scans.

1. Click **Trust > Compliance > Checklist**.

A list of existing checklists are displayed.

2. Click the checklist that you want to update.

Checklist details, rules and their statuses, and imported scans are listed.

3. Click the **Imported Scans** tab.

Scans available for importing are listed.

4. Click **Add**.

The Update Checklist page is displayed.

5. Select the scan that you want to import and click **Next**.

Rules from the selected scan are displayed.

6. Review the rules whose status is **Open**, and change the status to **Not Applicable** if the rule doesn't apply to the target, or **Resolved** if you have manually resolved the rule.

7. Click **Next** and then click **Save Checklist**.

The updated checklist is listed on the Compliance Checklists page.

Trust Plans Overview

A trust plan or a score plan defines how to calculate a trust score for a network entity. It comprises a set of trust factors for each factor category.

A trust plan also defines the following:

- Contribution values for each of the factors in the variable and reputational categories, describing the significance of the factor relative to other factors in the same category.

- Contribution factors for the variable and reputational categories defining the percentages that each category contributes toward the total trust score.

A trust score plan is applied on a network entity by

- Calculating the trust score based on the factors defined by the plan and the latest values of those factors for the network entity.
- Generating and persisting a trust score result.

Contribution values are associated with the trust factors in a trust score plan and are used to define the contribution of a factor to the calculation of the trust score. How the contribution values are used depends on the type of trust factor they are associated with.

A trust factor has an implied maximum contribution and an actual contribution. The percentage score for a category is determined to be **(the sum of the actual contributions for each of its factors/sum of the maximum contributions) * 100**.

The overall percentage score is derived from the variable and reputational percentages, adjusted according to these category contributions.

The compliance benchmark documents provided in this initial release are based on the documents that are published by the Center for Internet Security (CIS).

The CIS Controls Implementation Groups (IGs) are divisions that are based on cybersecurity characteristics. Each group identifies a subset that is assessed to be reasonable for an organization with a similar risk profile and resources to aim to implement. These IGs represent a cross-section of the CIS Controls that are customized to different types of businesses. Each IG builds atop the previous. For example, IG2 includes IG1, and IG3 includes all the Controls in IG1 and IG2.

CIS recommends that businesses prioritize their standardization of the Controls by inheriting from the IGs. Businesses should implement Controls in IG1, followed by IG2, and later IG3. The Controls contained within IG1 are critical to success. Support for IG1 should be considered among the first things to be done as part of a cybersecurity program. CIS describes IG1 as **Cyber Hygiene** – the essential protections that must be enforced to defend against common attacks.

In the case of Compliance, the IGs are used as a guidance to allocate an appropriate contribution against each trust factor, or a rule in the benchmark document.

Figure 46: Trust Score Generation Process



RELATED DOCUMENTATION

| [Trust Score Overview](#) | 621

About the Network Score Formula Page

IN THIS SECTION

- [Tasks You Can Perform](#) | 620
- [Field Descriptions](#) | 620

To access this page, click **Trust > General > Trust Plans**.

Routing Director generates a trust score for each target based on a trust score plan. A trust score plan defines how to calculate a trust score for a target. It comprises a set of trust factors for the prerequisite, variable, and reputational factors. The score plan also defines the contribution values for each factor in the variable and reputational factor categories. You can view the trust plans that are applied to the network from the Network Score Formula page.

Tasks You Can Perform

You can perform the following tasks from this page:

- View predefined trust plans available in Routing Director. You can also view details of the rule groups and the prerequisite, variable, and reputational factors that contribute to the trust score. Click a trust plan to view details of the rules defined in the plan.
- You can also perform the following tasks from this page:
 - Sort, resize, or re-arrange columns in a table (grid).
 - Show or hide columns in the table or reset page preferences, using the vertical ellipsis menu.
 - Search by using keywords—Click the search icon (magnifying glass), enter the search term in the text box, and press Enter. The search results are displayed on the same page.
 - Filter the data displayed in the table—Click the filter icon (funnel) and select whether you want to show or hide advanced filters. You can then add or remove filter criteria, save criteria as a filter, apply or clear filters, and so on. The filtered results are displayed on the same page.
 - Set the number of rows to be displayed in the table—Set the number of rows (15, 30, 50, or 100) to be displayed in the Network Implementation Plan table in the **Display** drop-down list provided at the bottom right-corner of the table.

Use the arrow heads to traverse through the pages or enter a page number to view the entries listed on that page.

Field Descriptions

[Table 146 on page 620](#) describes the fields in the Network Score Formula page.

Table 146: Field Descriptions on the Network Score Formula Page

Field	Description
ID	Unique identifier for the score plan.

Table 146: Field Descriptions on the Network Score Formula Page *(Continued)*

Field	Description
Name	Name of the score plan.
Version	Version of the score plan.
Variable Weighting	Percentage of weighting value assigned to variable factors in the score plan. Variable factors include a target's configuration, version, activated features, and so on
Prerequisite Factors	Number of prerequisite factors defined in the score plan.
Variable Factors	Number of variable factors defined in the score plan.
Reputational Factors	Number of reputational factors defined in the score plan.
Last Updated	Date when the score plan was last updated.

RELATED DOCUMENTATION

[Trust Plans Overview](#) | 617

[Trust Score Overview](#) | 621

Trust Score Overview

A trust score in Routing Director represents a level of trust in a network entity, expressed as a percentage, with 100 percent representing complete trust. A trust score is calculated based on the values of a list of factors, with contributions that reflect the relative significance of specific factors or groupings of factors.

A trust factor is any value of a network entity that can contribute to a trust score. The factors include metadata, such as identity, name, and description.

A factor category identifies the type of trust factor, such as prerequisite, variable, or reputational.

- **Prerequisites**—Conditions that a network target must meet to receive a nonzero trust score.
- **Variable Contributions**—Factors that provide a weighted trust contribution. You can assign weights based on nodes' characteristics, such as configurations, versions, and active features. Variable trust changes could be due to discrete events resulting in step changes, for example, activating a feature or upgrading a node.
- **Reputational Contributions**—Incremental trust contributions earned over time. It is a cumulative function of specified historical events, for example, number of times a node was reconfigured or spontaneous reboots.

A score plan defines how to calculate a trust score for a network entity. It comprises a set of trust factors for each factor category.

The score plan also defines:

- Contribution values for each of the factors in the variable and reputational categories, describing the significance of the factor relative to other factors in the same category
- A weighting for the variable and reputational categories defining the percentages of the total trust score that each category contributes

A score plan is applied on a network entity by:

- Calculating the trust score based on the factors defined by the plan and the latest values of those factors for the network entity
- Generating and persisting a trust score result.

Contribution values are associated with the trust factors in a score plan. These values are used to define the contribution of a factor to the calculation of the trust score. How contribution values are used depends on the type of trust factor with which those values are associated.

A trust factor has an implied maximum contribution and an actual contribution. The percentage score for a category, for example, the variable contribution is determined to be **(the sum of the actual contributions for each of its factors/sum of the maximum contributions) * 100**.

The overall percentage score is derived from the variable and reputational percentages, adjusted according to the category contribution weighting.

RELATED DOCUMENTATION

[Trust Plans Overview | 617](#)

[Trust and Compliance Overview | 603](#)

Manage Compliance Scans

IN THIS CHAPTER

- [Compliance Scans Overview | 623](#)
- [About the Compliance Page | 624](#)
- [Perform Custom Compliance Scans | 626](#)
- [Analyze Scan Results | 630](#)

Compliance Scans Overview

IN THIS SECTION

- [Labels | 624](#)

A Security Content Automation Protocol (SCAP) scan is the process for using known standards to run vulnerability and compliance scans. An SCAP scan allows the user to evaluate and secure their targets and networks.

An SCAP scan compares the system you are scanning to a baseline benchmarks document. The output of a SCAP scan is a SCAP results document.

The Compliance page displays a list of scans already run on the network. You can click a scan name to view details of the network targets scanned along with compliance scores in the range 0 – 100 for each target. A score of zero (0) indicates that the target device did not meet the compliance prerequisites to assign a valid score. A compliance score of 100 indicates that the target device is fully compliant.

Labels

In Routing Director, labels are key-value pairs that are attached to objects, such as a compliance scan, device registration, trust score plan, and so on. Each object can have a set of key-value labels defined. Each key must be unique for a given object.

Labels are used only for identifying objects, and can be used to organize subsets of objects. Labels can be attached to objects at the time of creation, and subsequently added and modified at any time through APIs.

RELATED DOCUMENTATION

[Perform Custom Compliance Scans | 626](#)

[Analyze Scan Results | 630](#)

About the Compliance Page

IN THIS SECTION

- [Tasks You Can Perform | 624](#)

To access this page, click **Trust > Compliance > Compliance Scan**.

The Compliance page displays a list of scans already run on the network. You can click a scan name to view details of the network targets scanned. The scan reports contain details such as the number of targets that are scanned, time taken for the scan, compliance score of the targets, and the rule results for the selected target. Additionally, you can view scan summary cards that display the number of targets that have a compliance score below a certain threshold, targets that are non-compliant, and targets that are fully compliant.

Tasks You Can Perform

You can perform the following tasks from this page:

- View scans that have been completed earlier.

- Run a custom scan. Click **Add** to initiate a custom scan. See ["Perform Custom Compliance Scans" on page 626](#).
- Search for compliance scans in which specific targets were scanned by using the entering the target name in the **Search By Target** field.
- You can also perform the following tasks from this page:
 - Sort, resize, or re-arrange columns in a table (grid).
 - Show or hide columns in the table or reset page preferences, using the vertical ellipsis menu.
 - Search by using keywords—Click the search icon (magnifying glass), enter the search term in the text box, and press Enter. The search results are displayed on the same page.
 - Filter the data displayed in the table—Click the filter icon (funnel) and select whether you want to show or hide advanced filters. You can then add or remove filter criteria, save criteria as a filter, apply or clear filters, and so on. The filtered results are displayed on the same page.
 - Set the number of rows to be displayed in the table—Set the number of rows (15, 30, 50, or 100) to be displayed in the Network Implementation Plan table in the **Display** drop-down list provided at the bottom right-corner of the table.

Use the arrow heads to traverse through the pages or enter a page number to view the entries listed on that page.

[Table 147 on page 625](#) describes the fields on the Scans page.

Table 147: Fields on the Compliance Page

Field	Description
Scan UUID	Unique identifier that Routing Director generates for the scan.
Scan Name	Auto-generated name for the scan.
Benchmark ID	Unique identifier that Routing Director generates for the benchmarks document used in the scan.
Benchmark Source	Source of the benchmarks document. For example, Center for Internet Security (CIS). Routing Director uses the benchmarks document that is defined by CIS.

Table 147: Fields on the Compliance Page (*Continued*)

Field	Description
Benchmark Name	Name of the benchmarks document used in the scan.
Benchmark Version	Version of the benchmarks document used in the scan.
Tailoring Name	Name of the tailorings document.
Tailoring Version	Version of the tailorings document.
Profile	Profile level selected for the scan.
Labels	Labels that you assign to a scan. As scan names are auto-generated, assigning a label helps you identify scans that are initiated by you.
Total Targets	Number of targets assessed during the scan.
Targets	Name of the target assessed during the scan. Hover over the target name to view the hostname, IP address, MAC address, and friendly name of the target.
Time Started	Date and the time when the scan was started.
Duration	Duration of the scan in milliseconds.
Status	Compliance status of the targets that were scanned.

Perform Custom Compliance Scans

Routing Director automatically runs scans to assess the targets in the network. While automatic scans check for compliance of all targets in the network, you can initiate custom scans to scan specified targets.

You can run a custom scan in the following ways:

- Adding details directly on the UI.

- Uploading a preconfigured JSON File.

To run a custom scan:

1. Click **Trust > **Compliance** > **Compliance Scan**.**

The Compliance page appears.

2. Click the add icon (+).

The Create Compliance Scan wizard appears.

3. Choose any one of the following options to configure the custom compliance scan:

- To enter details directly on the UI, perform the steps that are described in ["Enter Details Directly" on page 627](#).



NOTE: If you choose to enter details directly on the UI, do not upload a JSON file.

- To upload a preconfigured JSON file, perform the steps that are described in ["Upload JSON file" on page 628](#).

Enter Details Directly

a. Enter a name for the scan in the **Scan Name field. For example, scan-1.**

The name can contain alphanumeric characters and some special characters [hyphen (-), underscore (_), period (.), and colon (:)] and cannot exceed 64 characters. If you do not enter a name, the scan name is auto-generated.

b. Enter a description for the scan in the **Scan Description text box.**

c. Click **Next to view the Select Benchmark section. The source, benchmarks document, and version are selected by default.**

d. Select a benchmark profile depending on the level of security of the scan to be performed.

A benchmarks document may have one or more profiles. The value **<default>** indicates that you haven't selected a security profile. For more information, see [Table 1 on page 629](#).

Click **Clear** to clear the information that you provided in the Select Benchmark section.

e. Click **Next to view the Tailor section.**

f. (Optional) Select a tailoring document and version from the Name and Version drop-down lists respectively.

Alternatively, click **Clear** to clear the information that you provided in the Tailor section.

g. Click **Next to view the Select Targets section.**

- h. On the Select Targets section, select one or more targets that you want to scan from the **Available Targets** box and click the > icon to move the targets to the **Selected Targets** box.
- i. Click **Next** to view the Labels section.
- j. (Optional) On the Labels section, click add (+) to define a key-value pair. Labels (keypair) help you identify scans that you initiated. You can use these labels to filter completed scans.
- k. Click **Next** to view the Review section.

You can review the information that you entered on the Create Compliance Scan wizard in this section. The three accordions in this section provide information about benchmarks details, tailoring details, and labels assigned to the specified targets. The accordions are:

- i. **Benchmark Details**—View the rules that are defined in the benchmarks document for the scan. See ["About the Compliance Benchmarks Page" on page 608](#).
- ii. **Tailoring Details**—View the selected tailoring values for the scan. See ["About the Compliance Tailoring Page" on page 610](#).
- iii. **Add Labels**—View the defined key-value pair for the scan. See ["Compliance Scans Overview" on page 624](#).

(Optional) To edit information in any section of the wizard, click **Back** to navigate to that section. After you edit information, click **Next** to proceed to the next section.

Upload JSON File

- a. Click **Browse** to upload a preconfigured custom compliance scan configuration file in the JSON format.



NOTE: Ensure that the JSON file contains all the details required for each field that you want to scan.

Alternatively, to remove the uploaded JSON file, click **Clear**.

- b. (Optional) Enter a name for the scan in the **Scan Name** field. For example, scan-1.

The name can contain alphanumeric characters and some special characters [hyphen (-), underscore (_), period (.), and colon (:)] and cannot exceed 64 characters. If you do not enter a name, the default scan name is auto-generated.

- c. (Optional) Enter a description for the scan in the **Scan Description** text box.
- d. Click **Next** to view the Review section of the Create Compliance Scan wizard.

The fields (see [Table 1 on page 629](#)) on the Create Compliance Scan wizard are populated with the information available in the JSON file that you uploaded.

You can review these details in this section. The three accordions in this section provide information about benchmarks details, tailoring details, and labels assigned to the specified targets. The accordions are:

- i. **Benchmark Details**—View the rules that are defined in the benchmarks document for the scan. See ["About the Compliance Benchmarks Page" on page 608](#).
- ii. **Tailoring Details**—View the selected tailoring values for the scan. See ["About the Compliance Tailoring Page" on page 610](#).
- iii. **Add Labels**—View the defined key-value pair for the scan. See ["Compliance Scans Overview" on page 624](#).

4. Do any of the following:

- Click **Cancel** to exit the wizard without saving changes.
- Click **Back** to go to the Labels page.
- Click **Export** to download the details in the JSON file format.

You can ["upload this file" on page 628](#) when you want to run a custom compliance scan with the same details available in the JSON file that you exported.

- Click **Scan** to initiate the scan.

The newly initiated scan is listed on the Compliance page. After the scan is completed, you can analyze the scan results. For more information, See ["Analyze Scan Results" on page 630](#).

Table 148: Fields on the Create Compliance Scan Wizard

Field	Description
Source	Select the organization that provides the benchmarks document. For example, Center for Internet Security (CIS).
Benchmark	Select the benchmarks document applied on the network.
Version	Select the version of the benchmarks document.

Table 148: Fields on the Create Compliance Scan Wizard *(Continued)*

Field	Description
Profile	Select a security profile. A typical benchmarks document has three recommended profiles: default, Level 1 and Level 2. While the profile Level 1 is the base recommendation that doesn't cause much performance impact, Level 2 is for environments that need stricter security enforcement. The default profile is applied if no profile is selected.

RELATED DOCUMENTATION

[Analyze Scan Results](#) | 630

Analyze Scan Results

You can analyze the scan results to identify target devices that do not comply with the rules in the benchmark and tailoring documents and take corrective action to improve the trust score.

To analyze scan results:

1. Click **Trust > Compliance > Compliance Scan**.
The Compliance page appears displaying completed scans.
2. Click the scan name to view more details of target devices that are non-compliant, such as the rules that failed.
3. Click a compliance score to view the list of rules that the target was evaluated against.
4. Click the **Details** icon next to the Rule ID that has a status **Fail** in the **Status** column.
A panel showing more details, such as the reason for failure and the recommended resolution is displayed on the right side.
5. Perform the recommended action and rerun the scan.
6. Verify that the status of the Rule ID is displayed as **Pass**.

Manage Vulnerabilities

IN THIS CHAPTER

- [Vulnerabilities Overview | 631](#)
- [About the SIRT Advisories Page | 632](#)
- [About the Proactive Bug Notifications Page | 635](#)

Vulnerabilities Overview

The Juniper Networks Security Incident Response Team (Juniper SIRT) constrains the publication of Juniper Security Advisories and Security Notices for non-urgent issues to a predefined quarterly schedule. The schedule is as follows: the second Wednesday of January, April, July, and October, covering all Juniper products.

In exceptional circumstances, the Juniper SIRT might publish an out-of-cycle Security Advisory or Security Notice. Examples include active malicious exploitation of a zero-day Juniper vulnerability or a multi-vendor issue in which all participating parties must publish simultaneously on a schedule negotiated by an external coordinating agency. Routing Director provides information only about those SIRT advisories that are already released and does not include newer SIRT advisories published out of cycle after the release of Routing Director.

The Juniper SIRT considers numerous criteria for determining whether an issue warrants SIRT attention and, if so, how a fix will be applied and to what range of products and software releases, and how and when the issue will be published. The Juniper SIRT uses the Common Vulnerability Scoring System (CVSS) to rank an issue as one factor in its evaluation.

For more information, see the [Common Vulnerability Scoring System \(CVSS\) and Juniper's Security Advisories](#) page.



NOTE: If a target type is affected by a SIRT advisory, it does not imply that the target instance in your network is also affected. You need to investigate further to determine whether the problem definition and matching criteria are relevant to your deployment.

Juniper SIRT investigates such incidents and provides a comprehensive analysis of the security exposure that is based on your hardware, installed software, and configuration.

The Vulnerabilities page lists all the SIRT advisories that Juniper Networks has published, the devices on the network affected by these advisories, and the common vulnerabilities and exposures (CVEs).

About the SIRT Advisories Page

IN THIS SECTION

- [Tasks You Can Perform | 632](#)
- [Field Description | 633](#)

To access this page, click **Trust > Vulnerabilities > Advisories**.

Routing Director regularly monitors the targets in the network for vulnerabilities and potential security risks, and generates alerts. These alerts contain details of the vulnerability, its potential impact, and recommendations for remediation. Network administrators can use these alerts to perform corrective actions.

According to severity, the advisories are classified as Critical, High, and Low. To view more information about the SIRT advisory, click the **Details** icon. You can view how the score in the CVSS Score column is arrived at and the CVE details.

By default, the Vulnerabilities page displays only the SIRT Advisories relevant to the products installed in the network.

Tasks You Can Perform

You can perform the following tasks from this page:

- View Juniper SIRT Advisories (JSA) which describe vulnerabilities in Juniper software, corresponding CVEs, and the proposed resolution. Use the search option to search for advisories by specific or a device model. Click the **Details** icon to view more information about an advisory. For more information, see "[Vulnerabilities Overview](#)" on page 631.
- You can also perform the following tasks from this page:

- Sort, resize, or re-arrange columns in a table (grid).
- Show or hide columns in the table or reset page preferences, using the vertical ellipsis menu.
- Search by using keywords—Click the search icon (magnifying glass), enter the search term in the text box, and press Enter. The search results are displayed on the same page.
- Filter the data displayed in the table—Click the filter icon (funnel) and select whether you want to show or hide advanced filters. You can then add or remove filter criteria, save criteria as a filter, apply or clear filters, and so on. The filtered results are displayed on the same page.
- Set the number of rows to be displayed in the table—Set the number of rows (15, 30, 50, or 100) to be displayed in the Network Implementation Plan table in the **Display** drop-down list provided at the bottom right-corner of the table.

Use the arrow heads to traverse through the pages or enter a page number to view the entries listed on that page.

Field Description

Table 149 on page 633 lists the fields on the Vulnerabilities page.

Table 149: Fields on the SIRT Advisories Page

Field	Description
ID	Unique identifier that Routing Director generates for the security advisory.
SIRT ID	Unique SIRT ID for the security advisory. The Juniper Networks Security Incident Response Team (Juniper SIRT) generates this ID for the security advisory.
Title	Title of the security advisory.
Date	Date on which the security advisory was first published.
Severity	Severity rating of the security advisory as Critical, High, or Low.
CVSS Score	Severity score of the advisory in the range 0-10.

Table 149: Fields on the SIRT Advisories Page *(Continued)*

Field	Description
Targets	Name of the target affected by the security advisory. Hover over the target name to view the hostname, IP address, MAC address, and friendly name of the target.
Total Targets	Number of targets affected by the security advisory.
Products Affected	Number of products affected by the security advisory. Click the value to see the details of products affected.
CVEs	Common Vulnerabilities and Exposures (CVE) IDs of the security advisory. For example, CVE-2023-1234 is for publicly known vulnerabilities.
Problem	Description of the problem.
Workaround	Workaround for the problem.
Models	Model of the target. For example, EX2200.
Model Count	Number of device models.
Target ID	Unique identifier that Routing Director generates to uniquely identify each target.
Acknowledged	Indicates whether you have acknowledged the issue.

About the Proactive Bug Notifications Page

IN THIS SECTION

- [Tasks You Can Perform | 635](#)
- [Fields on the Proactive Bug Notifications Page | 636](#)

To access this page, click **Trust > Vulnerabilities > PBNs**.

Routing Director alerts you in advance about the problems in the devices in the network. The proactive bug notifications (PBN) feature in Routing Director enables you to identify potential bugs in the devices so that you can plan software upgrades. The total number of bugs that affect the devices in the network is categorized according to their severity and are displayed in the insights bar at the top of the page. The information displayed in the insights bar makes it easy to identify devices that need a software upgrade.

Tasks You Can Perform

- **Acknowledge a bug**—Acknowledge a bug to indicate that the bug has come to your notice. Select a bug and click **Acknowledge** from the **More** drop-down list. Acknowledging a bug doesn't mean that the bug is resolved. An acknowledged bug is excluded from the total bug count in the insights bar. You can click **Hide Acknowledged** to display only critical bugs or bugs that need to be reviewed.
- **Unacknowledge a bug**—Unacknowledge a bug to make sure that the network administrator can continue tracking the bug. Select a bug and click **Unacknowledge** from the **More** drop-down list.
- **View bug details**—View detailed information about a bug, such as the description, problem level, PR score that indicate the relative criticality of the bug, release in which the bug is resolved, and so on.
- **Filter the data displayed in the table**—Click the filter icon (funnel) and select whether you want to show or hide advanced filters. You can then add or remove filter criteria, save criteria as a filter, apply or clear filters, and so on. The filtered results are displayed on the same page. By default, only those PBNs that are unacknowledged and affect at least one device are displayed.
- You can also perform the following tasks from this page:
 - Sort, resize, or re-arrange columns in a table (grid).
 - Show or hide columns in the table or reset page preferences, using the vertical ellipsis menu.
 - Search by using keywords—Click the search icon (magnifying glass), enter the search term in the text box, and press Enter. The search results are displayed on the same page.

- Filter the data displayed in the table—Click the filter icon (funnel) and select whether you want to show or hide advanced filters. You can then add or remove filter criteria, save criteria as a filter, apply or clear filters, and so on. The filtered results are displayed on the same page.
- Set the number of rows to be displayed in the table—Set the number of rows (15, 30, 50, or 100) to be displayed in the Network Implementation Plan table in the **Display** drop-down list provided at the bottom right-corner of the table.

Use the arrow heads to traverse through the pages or enter a page number to view the entries listed on that page.

Fields on the Proactive Bug Notifications Page

Table 150 on page 636 lists the fields on the Proactive Bug Notifications page.

Table 150: Fields on the Proactive Bug Notifications Page

Field	Description
ID	Unique ID that Routing Director generates for the bug.
PR Number	Unique ID generated by the bug tracking system to track the bug.
Headline	Brief description of the bug.
Customer Risk	Indicates the severity of the bug and its impact on the customer network.
Problem Level	Indicates the level of impact on the network.
PR Score	System generated score that indicates the impact of the bug. A higher score indicates that the bug is of higher severity.
Platform Category	Device series or models that the bug affects.
Description	Detailed description of the bug.
Trigger	Scenario in which the problem occurs.

Table 150: Fields on the Proactive Bug Notifications Page (*Continued*)

Field	Description
Fixed In	Releases in which the bug is resolved. Helps you determine the release to which the device software needs to be upgraded to resolve the issue.
Targets	List of devices in the network that are affected by the bug.
Total Targets	Total number of devices that are affected by the bug.
Acknowledged	Indicates whether you as an administrator have noticed the issue. Acknowledging a bug doesn't mean that the bug is resolved.

Monitor Integrity

IN THIS CHAPTER

- [Integrity of the Hardware and Software on the Network | 638](#)
- [About the Software End of Life Page | 639](#)
- [About the Hardware End of Life Page | 641](#)

Integrity of the Hardware and Software on the Network

Routing Director periodically notifies you about the integrity of the devices and software running on the devices in your network.

Routing Director maintains a database of the latest Juniper Networks hardware and software releases. In addition, Routing Director periodically collects information about the devices on the network and the version of software running on the devices. It then compares the collected information against the information that is maintained in the database. The information is collected to ascertain whether the devices on the network and the software running on these devices are in line with the vendor's recommendation. Routing Director notifies you in advance when a device or the software running on the device nears its end of life (EOL).

RELATED DOCUMENTATION

- [About the Hardware End of Life Page | 641](#)
- [About the Software End of Life Page | 639](#)

About the Software End of Life Page

IN THIS SECTION

- [Tasks You Can Perform | 639](#)
- [Field Description | 640](#)

To access this page, click **Trust > Integrity > Software EOL**.

The Software End of Life page helps you monitor the integrity of the OS running on the devices in the network and keeps the OS up to date with the latest supported releases.

Routing Director automatically tracks the versions of software running on the targets in the network. It provides information about a device whose OS is nearing its EOL date or is already past its EOL date. The graphical timeline provides the OS software EOL information at a glance.

The page provides a red, yellow, or green icon (next to the OS version) depending on the EOL date.

- A red icon indicates that the software has crossed the EOL date.
- A yellow icon indicates that the software is nearing its EOL date.
- A green icon indicates that the software's EOL date is not in the immediate future.

You can view the software support information, such as the release date, EOL date, and the date after which the software will not be supported, by clicking the software version on the timeline. You can also see a list of devices running each software version. Click the EOL date to view a list of devices that will reach EOL on that date.

Tasks You Can Perform

You can perform the following tasks from this page:

- View devices on the network whose OS software is nearing EOL or have already crossed the EOL date. The graphical timeline provides the OS software EOL information at a glance. Click the EOL date or the OS version on the timeline to view detailed information.
- You can also perform the following tasks from this page:
 - Sort, resize, or re-arrange columns in a table (grid).
 - Show or hide columns in the table or reset page preferences, using the vertical ellipsis menu.

- Search by using keywords—Click the search icon (magnifying glass), enter the search term in the text box, and press Enter. The search results are displayed on the same page.
- Filter the data displayed in the table—Click the filter icon (funnel) and select whether you want to show or hide advanced filters. You can then add or remove filter criteria, save criteria as a filter, apply or clear filters, and so on. The filtered results are displayed on the same page.
- Set the number of rows to be displayed in the table—Set the number of rows (15, 30, 50, or 100) to be displayed in the Network Implementation Plan table in the **Display** drop-down list provided at the bottom right-corner of the table.

Use the arrow heads to traverse through the pages or enter a page number to view the entries listed on that page.

Field Description

Table 151 on page 640 lists the fields on the Software End of Life page.

Table 151: Fields on the Software End of Life Page

Field	Description
Target ID	Identifier that Routing Director generates to uniquely identify each target.
Target	Name of the target. Hover over target name to view the hostname and IP address of the target.
Hostname	Hostname of the target.
IP Address	IP address of the target.
Friendly Name	Name used to easily identify the device.
Manufacturer	Device manufacturer name. For example, Juniper Networks.
Model	Device model name.
OS Version	Version of OS running on the target.
OS Name	Name of the OS. For example, Junos OS.

Table 151: Fields on the Software End of Life Page *(Continued)*

Field	Description
First Release Shipping	Date on which the target was first released.
End Of Life	EOL date of the target.
End of Engineering	Displays the date on which the engineering support for the specified software version running on the target ends. After this date, engineering support such as updates, bug fixes, and patches is no longer provided.
Discovery Status	Indicates whether the target is connected or disconnected to Routing Director.

About the Hardware End of Life Page

IN THIS SECTION

- [Tasks You Can Perform | 641](#)

To access this page, click **Trust > Integrity > Hardware EOL**.

Routing Director automatically tracks the end of life (EOL) information of managed devices in the network and their individual components whose EOL dates have been announced. Routing Director builds an EOL hardware inventory of all the hardware components in the network by matching the discovered devices with the information available about their components in the EOL database. During this process, if it identifies affected components, they are flagged on the Hardware End of Life page.

Tasks You Can Perform

You can perform the following tasks from this page:

- View details of managed devices and their individual components, whose EOL dates have been announced. To view more information about a device or a component, click the **Details** icon that is displayed when you hover over the SKU name or select a SKU and click **More > Details**. The SKU details pane displays the component details, support information, targets on which these hardware components are present, and recommended replacement device or components.
- You can also perform the following tasks from this page:
 - Sort, resize, or re-arrange columns in a table (grid).
 - Show or hide columns in the table or reset page preferences, using the vertical ellipsis menu.
 - Search by using keywords—Click the search icon (magnifying glass), enter the search term in the text box, and press Enter. The search results are displayed on the same page.
 - Filter the data displayed in the table—Click the filter icon (funnel) and select whether you want to show or hide advanced filters. You can then add or remove filter criteria, save criteria as a filter, apply or clear filters, and so on. The filtered results are displayed on the same page.
 - Set the number of rows to be displayed in the table—Set the number of rows (15, 30, 50, or 100) to be displayed in the Network Implementation Plan table in the **Display** drop-down list provided at the bottom right-corner of the table.

Use the arrow heads to traverse through the pages or enter a page number to view the entries listed on that page.

Table 152: Field Descriptions

Field	Description
SKU ID	Unique identifier generated for the SKU.
SKU	Displays the SKU name for the component. To view more information about a component, click the Details icon that is displayed when you hover over the SKU. The SKU Details pane displays the component details, support information, and the targets on which these hardware components are present.
Total Targets	The number of devices on which the SKU is used.
Targets	Displays the target associated with the SKU.
PSN Number	Displays the product support notification number for the SKU.

Table 152: Field Descriptions *(Continued)*

Field	Description
Description	Displays a description of the target.
Announced	Displays the date on which the EOL was announced.
Last Order	Displays the date on which the last order for the SKU can be made.
Last Date to Convert to Warranty	Displays the last date to convert the SKU's standard warranty into an extended support contract. You must complete this conversion before the standard warranty expires.
Same Day Support Discontinued	Displays the date on which the same-day support for the target discontinues.
Next Day Support Discontinued	Displays the date on which the next-day support for the target discontinues.
End of Support	Displays the date on which the support agreement expires.
End of Engineering	Displays the date on which the engineering support for the specified software version of the SKU ends. After this date, engineering support including maintenance releases and software updates for the specified SKU is no longer available.
End of Engineering HW	Displays the date on which the engineering support for the hardware ends. After this date, component support and associated firmware updates are no longer available.
Date Created	Displays the date on which the EOL for the SKU was created.

Table 152: Field Descriptions *(Continued)*

Field	Description
Date Updated	Displays the date on which the EOL for the SKU was updated.
Replacements	Displays recommended replacement component for the component that has reached its end of life.
Target ID	Unique identifier that Routing Director generates to uniquely identify each target.



Service Orchestration

- [Introduction | 646](#)
 - [View Service Design Catalog | 651](#)
 - [Manage Customers | 664](#)
 - [Add Resources for Network Services | 669](#)
 - [Provision VPN Services | 773](#)
 - [Manage Service Instances | 828](#)
 - [Monitor Service Order Execution Workflows | 874](#)
-

Introduction

IN THIS CHAPTER

- [Service Orchestration Overview | 646](#)
- [Service Provisioning Workflow | 648](#)

Service Orchestration Overview

Service orchestration is the process of designing, configuring, validating, deploying, and monitoring a network service; for example, Layer 3 VPN (L3VPN) service. From the service orchestration perspective, a network service is any point-to-point, point-to-multipoint, or multipoint-to-multipoint connection. Routing Director provides an automated framework that manages the entire life cycle of a network service, right from design to deprovisioning the service from the network.

Service orchestration in Routing Director is model-driven and intent-based. That is, a network administrator can use service designs available in the service catalog to provision a service. A service design defines how to create the service. You can create service instances using a service design. A service instance defines the sites, devices, connections, and other parameters of the service. To provision the service instance, a service order is created, which starts the service provisioning workflow. The workflow allocates the resources (such as access circuit, route target, route distinguisher, and so on) that satisfy the requirements of the service, produces a network-centric view of the service, and deploys the service in the network.

You can provision network services using the Routing Director GUI, or programmatically by using REST API. Superusers can also provision services using the service orchestration CMGD CLI.

Routing Director provides the following service and resource designs to onboard devices to the network and to provision network services on the onboarded devices:

- Infrastructure service design to onboard and deploy devices to the network.
- L3VPN service design to provision and monitor L3VPN services in the network.
- E-LAN EVPN CSM or E-Line EVPN VPWS CSM service design to provision and monitor EVPN services in the network.

- E-Line L2circuit NSM service design to provision and monitor L2 circuit services in the network.
- Routing and Layer 3 resource designs to configure network resource pools for infrastructure services.
- Layer 2 resource design to configure network resource pools for EVPN services.
- Topology resource design to configure topology resource pools for L3VPN, EVPN, and EVPN-VPWS services.
- VPN-resources resource design to configure VPN resource pools for L3VPN, EVPN, and L2 circuit services.

The service and resource designs are preinstalled when you install Routing Director.

Service Orchestration in Routing Director consists of the following elements:

- **Service Catalog**—A service catalog is a collection of predefined service designs that are available for the organization. Service designs define how to create a service.
- **Service Designs**—Service orchestration in Routing Director is model-driven, that is, service orchestration is carried out through predefined service designs, which contain the specifications for a type of service.
- **Service Instance**—A service instance is a specific instance of a service design created for a customer. For example, a named instance of an L3VPN service created for a specific customer. Service instances define the sites, devices, connections, and other parameters of the service.
- **Resource Designs**—A resource design is a special type of service design whose only purpose is to produce resources (such as access circuits, route targets, and so on) which are consumed by network services. Resource designs generally do not affect the configuration of the network devices.
- **Resource Instance**—A resource instance is a specific instance of a resource design that defines the configurations for a resource pool required to provision a service in the network.
- **Placement**—Process of allocating network resources for provisioning a service instance.
- **Service Order**—Process of creating, modifying, or deleting a service instance. A service instance is implemented through a service order. The provisioning status of the service instance can be accessed through the service order.
- **Workflow**—A workflow is a sequence of tasks that Routing Director executes to provision a service. You can monitor workflows and access logs that help you troubleshoot failed service orders.

To view a list of devices that support service provisioning, see [Supported Junos OS Releases, Devices, and Browsers](#).



NOTE: You can provision services on greenfield and brownfield devices.

RELATED DOCUMENTATION

- [Service Provisioning Workflow | 648](#)
- [About the Service Designs Page | 653](#)
- [About the Service Instances Page | 829](#)
- [About the Resource Instances Page | 672](#)
- [About the Service Orders Page | 874](#)
- [About the Service Orchestration cMGD CLI | 1576](#)

Service Provisioning Workflow

Service orchestration automates provisioning of Layer 3 VPN (L3VPN), Layer 2 Ethernet VPN (EVPN), Layer 2 EVPN Virtual Private Wire Service (VPWS), and Layer 2 circuit (L2 circuit) services and monitoring of the health and quality of a service after it is provisioned. Routing Director provides predefined service designs that codify the guidelines for executing the service provisioning and monitoring workflows for the services.



NOTE:

- LDP signaling without BGP autodiscovery is the only variant of L2 circuit that is supported in this release.
- Redundant pseudowires for L2 circuit are not supported in this release.

The service provisioning sequence of activities includes tasks such as configuring network resources, defining service elements and initiating service orders for service provisioning, monitoring service order execution states, and troubleshooting issues by monitoring task logs. A network administrator or a superuser provisions a service by using the Routing Director GUI, REST API, or the service orchestration cMGD CLI.

To provision a service:

1. Add customer names in the Customer Inventory page. See ["About the Customer Inventory Page" on page 664](#).

2. View the installed service designs with version numbers on the Service Designs page. Network administrators and observers have read-only access to this page. Superusers can manage the service design catalog and perform tasks like installing the latest available version of a service design and setting the default service design version. See ["About the Service Designs Page" on page 653](#) and ["Manage Service Design Versions" on page 657](#).

Superusers with root user privileges can manage service designs by using the service orchestration cMGD CLI. See ["About the Service Orchestration cMGD CLI" on page 1576](#).

3. Configure and upload network resources for the services by using the Routing Director GUI, REST API, or the service orchestration cMGD CLI. See ["Add Network Resources for Provisioning Services" on page 670](#), ["About the Resource Instances Page" on page 672](#), and ["request network resources load" on page 1631](#).
4. Create or modify service instances by entering data in the corresponding fields in the Routing Director GUI or by uploading preconfigured JSON files:
 - To create an L3VPN service instance, see ["Add an L3VPN Service Instance" on page 773](#).
 - To create an EVPN or EVPN-VPWS service instance, see ["Add an EVPN or EVPN-VPWS Service Instance" on page 791](#).
 - To create an L2 Circuit service instance, see ["Add an L2 Circuit Service Instance" on page 810](#).

To modify a service instance, see ["Modify a Service Instance" on page 822](#).

When you create or modify a service instance, a create or modify service order is generated for these operations.

5. Provision the service instance by clicking the **Provision** button on the Service Instances page. Provisioning a service instance generates a service order to activate the create or modify workflows.

You can also delete a service instance by clicking **More > Deprovision** on this page to generate a delete service order to activate the delete workflow. See ["About the Service Instances Page" on page 829](#).
6. View details of the service order and the execution status of the service order on the Service Orders page. See ["About the Service Orders Page" on page 874](#).
7. Monitor the workflow run details for each service order from the Workflows page (**Orchestration > Monitoring > Workflows**). By viewing detailed task logs, you can troubleshoot issues when a workflow run fails. See ["About the Workflows Page" on page 880](#).

As a superuser with root user privileges, you can perform service provisioning tasks by using the service orchestration cMGD CLI. For more information, see ["About the Service Orchestration cMGD CLI" on page 1576](#).

Before provisioning a service in the network, Routing Director validates the service configuration on all the devices on which the service is provisioned. Routing Director automatically configures passive device monitoring rules for all services and active data plane tests for L3VPN services to monitor the health and quality of the service after provisioning. You can view the monitoring data generated for each service instance. See "[View Service Instance Details](#)" on page 836.

RELATED DOCUMENTATION

About the Customer Inventory Page	 664
About the Service Designs Page	 653
About the Service Instances Page	 829
Add Network Resources for Provisioning Services	 670
About the Resource Instances Page	 672
Add an L3VPN Service Instance	 773
Modify a Service Instance	 822
About the Service Orders Page	 874
About the Workflows Page	 880
View Service Instance Details	 836
About the Service Orchestration cMGD CLI	 1576

View Service Design Catalog

IN THIS CHAPTER

- [Service Design Overview | 651](#)
- [About the Service Designs Page | 653](#)
- [Manage Service Design Versions | 657](#)
- [Upload a Customized Service Design | 659](#)

Service Design Overview

Service orchestration in Routing Director automates provisioning of a service in the network and monitoring the service after provisioning. Service designs codify the guidelines to onboard devices, configure network resources for a service, provision a service, and monitor the service after provisioning. A service design includes the following:

- A service model that defines the service structure and elements. Written in the YANG language, a service model includes a customer service model (CSM) and a network service model (NSM). The CSM is intent-based and describes a service from the customer's point of view. The NSM is implementation-based and describes a service from the network's point of view.
- Rules to validate whether a service instance that is configured based on a service design conforms to the service model in the service design. Validation rules are usually embedded in the CSM or NSM.
- Rules for allocation of resources (placement) to implement a service instance. Placement rules are jq filters that are encapsulated in a YAML file.
- Translation templates that describe how to transform service instances to service configurations for deployment on target devices. Translation templates are Jinja2 files that differ for different devices, vendor OSs, or software versions.
- Monitoring rules to monitor the service after provisioning by collecting relevant information from all devices on which the service is provisioned.
- Active data plane tests to test a service after provisioning it in the network.

- Active data plane tests to continually measure the quality of a provisioned service.

Routing Director provides the following service and resource designs for onboarding devices, adding network resources, and provisioning L3VPN, EVPN, EVPN-VPWS and L2 circuit services:

- Infrastructure service design to onboard and deploy devices to the network.
- L3VPN service design to provision and monitor L3VPN services in the network.
- E-LAN EVPN CSM service design to provision and monitor EVPN services in the network.
- E-Line EVPN VPWS CSM service design to provision and monitor EVPN-VPWS services in the network.
- E-Line L2circuit NSM service design to provision and monitor L2 circuit services in the network.
- Routing and L3 resource designs to configure network resource pools for infrastructure services.
- L2 resource design to configure network resource pools for EVPN services.
- Topology resource design to configure topology resource pools (site and PE details) for L3VPN, EVPN, and EVPN-VPWS services.
- VPN-resources resource design to configure VPN resource pools (route targets, route distinguishers, and virtual circuits) for the L3VPN, EVPN, and L2 circuit services.

Routing Director supports up to three concurrent versions of each service and resource design. The service and resource designs are automatically installed when you install Routing Director.

While the service and resource design catalog is automatically updated when you install Routing Director, service and resource designs for a given organization are not automatically upgraded. The upgrades are managed by the organization administrator. To upgrade a service design used by a service instance, you must explicitly upgrade the version using the GUI.

The Service Designs page (**Orchestration > Service > Service Catalog**) lists all the service designs and the available versions of each service design.

A root user can manage the service design catalog by using the service orchestration cMGD CLI. See ["About the Service Orchestration cMGD CLI" on page 1576](#).

RELATED DOCUMENTATION

[About the Service Designs Page | 653](#)

[Device Onboarding Overview | 171](#)

[Network Implementation Plan Overview | 230](#)

About the Service Designs Page

IN THIS SECTION

- [Tasks You Can Perform | 653](#)
- [Field Descriptions | 654](#)

The Service Designs page is an inventory of service designs that Routing Director provides for provisioning services. The page lists service designs used to provision L3VPN, EVPN, EVPN-VPWS, or L2 circuit services in your network. The service designs are listed by name, service type, version number, number of active instances, and so on. Routing Director allows you to install up to three concurrent versions of a service design in a network. One of the three versions is set as the default version.



NOTE: You must be a superuser to manage the service design catalog. Network administrators and observers have read-only access to the inventory.

To access the Service Designs page, click **Orchestration** > **Service Catalog** on the navigation menu.

Tasks You Can Perform

You can perform the following tasks on the Service Designs page:

- View service design details—To view details of a service design, do any of the following:
 - Hover over the service design name and click the **Details** icon that appears next to the design name.
 - Select the design and click **More** > **Show Detail**.

The *service-design-name* pane appears on the right side of the Service Catalog page. On the *service-design-name* pane, you can view general information about the service design and details of other service designs that this service design depends on. See [Table 154 on page 655](#) and [Table 155 on page 656](#).

- View the service instances associated with a service design—Click the *service-design-name* hyperlink. You are directed to the Service Instances page, where the service design name is used as a filter to display the service instances associated with the design.
- Set the default version of a service design. See ["Set Service Design Default Version" on page 657](#).

- Install the latest available version of a service design. See ["Install Latest Available Service Design Version" on page 657](#).
- Delete a service design.

Routing Director supports up to three concurrent versions of a service design. To ensure that there are no more than three versions of a service design at a given time and to install the latest available version, a superuser can delete a specific or an unused version of a service design for an organization.

To delete a service design:

1. Select a service design and click **More > Delete Service Design**.

The Delete Service Design confirmation page appears.

2. Click **Yes** to delete.

The service design is deleted and a confirmation message appears. You are returned to the Service Designs page. The deleted service design version is listed in the **Pending Version** column.



NOTE: You cannot delete a service design if there are active service instances associated with the service type.

- Sort, resize, or re-arrange columns in a table (grid).



NOTE: Sort functionality is not available for the Version, Pending Version, and # Active Instances columns.

- Refresh data displayed in the Service Designs table—Click the Refresh icon at the bottom-left corner of the Service Designs table to refresh the data displayed in the table.

Field Descriptions

[Table 153 on page 654](#) lists the fields on the Service Designs page.

Table 153: Fields on the Service Designs Page

Field	Description
Name	Name of the service design.

Table 153: Fields on the Service Designs Page (Continued)

Field	Description
Service Type	The type of service that the service design provides guidelines for: • Onboard—Service type for adding network infrastructure; for example, onboarding devices to Routing Director. • Network resource—Service type for adding routing, topology, Layer 3, Layer 2, and VPN resource pools to Routing Director. • L3VPN—Service type for provisioning L3VPN service in a network. • L2VPN—Service type for provisioning EVPN, EVPN-VPWS, and L2 circuit services in a network.
Version	Version number of the service design.
Pending Version	The latest available version of the service design pending installation.
Created time	Date and time when the service design was created.
Created by	Name of the user who created the service design.
# Active Instances	Number of active service instances associated with the design.
Description	Brief description about the service design, if any.

Table 154: Fields in the General Section of *service-design-name* Pane

Field	Description
Name	Name of the service design.

Table 154: Fields in the General Section of *service-design-name* Pane (Continued)

Field	Description
Service Type	Type of service for which the service design provides guidelines. • Onboard—Service type for adding network infrastructure; for example, onboarding devices to Routing Director. • Network resource—Service type for adding routing, topology, Layer 3, Layer 2, and VPN resource pools to Routing Director. • L3VPN—Service type for provisioning L3VPN service in a network. • L2VPN—Service type for provisioning EVPN, EVPN-VPWS, and L2 circuit services in a network.
Version	Version of the service design.
Release Notes	Summary of service design enhancements and bug fixes. Hover over the value of this field to view the summary of enhancements.

Table 155: Fields in the Dependencies Table

Field	Description
Name	Name of service design on which the current service design depends. For example, topology resource design on which the L3VPN service design depends.
Service Type	Service type offered by the service design on which the current service design depends. For example, network resource service type for topology resource design on which the L3VPN service design depends.
Version	Version of the service design on which the current service design depends.

RELATED DOCUMENTATION

[Service Design Overview](#) | 651

Manage Service Design Versions

IN THIS SECTION

- [Set Service Design Default Version | 657](#)
- [Install Latest Available Service Design Version | 657](#)
- [Upgrade Service Design Version for a Service Instance | 658](#)

As a superuser, you can manage the service design catalog and perform tasks such as setting the default version of a service design and installing the latest available version of a service design that Routing Director provides. You can perform these tasks by using the Routing Director GUI or the service orchestration cMGD CLI. Superusers with root user privilege can also uninstall unused versions of a service design by using the service orchestration cMGD CLI. See ["request service design uninstall" on page 1644](#).

Set Service Design Default Version

The Service Designs page lists up to three concurrent versions of a service design. One of the three versions is set as the default version and is marked as **(Default)** in the Service Designs table, next to the version number. Any service instance you create, modify, or delete, is associated with the default service design.

To set the default version of a service design by using the Routing Director GUI:

1. Click **Orchestration > Service Catalog**.
The Service Designs page appears.
2. Select the service design version you want to set as the default version.



NOTE: You must select a version that is not listed as **(Default)** in the table.

3. Click the **Set Default Version** button and click **Yes** to confirm.
You see a message that the default version is successfully changed. The table automatically refreshes to display the updated service design default version.

To set the default version of a service design by using the service orchestration cMGD CLI, see ["set service design default version" on page 1584](#).

Install Latest Available Service Design Version

Routing Director supports up to three concurrent versions of a service or resource design. The Service Designs page lists the latest available version of a service design that is pending installation. You can

install the latest available version by using the Routing Director GUI or the service orchestration cMGD CLI.

To install the latest service design version by using the GUI:

1. Click **Orchestration > Service Catalog**.

The Service Designs page appears.

2. Select the service design for which you want to install the latest available version.

3. Click **More > Install Pending Version**.

Alternatively, right-click the version number and click **Install Pending Version**.

4. On the Install Version dialog box, click **Yes** to confirm.



NOTE:

- The latest version you install is set as the default version for the service design. To change the default version, see ["Set Service Design Default Version" on page 657](#).
- Routing Director supports only three concurrent versions of a service design. To install the latest available version, superusers with root user privilege can uninstall an unused version of a service design by using the service orchestration cMGD CLI. See ["request service design uninstall" on page 1644](#).

To install the latest available version of a service design by using the service orchestration cMGD CLI, see ["request service design install" on page 1643](#).

Upgrade Service Design Version for a Service Instance

Any service instance you create, modify, or delete is associated with the default service design version. You can upgrade a service design version associated with a service instance by using the Routing Director GUI.

To upgrade a service design version for a service instance:

1. Click **Orchestration > Instances**.

The Service Instances page appears.

2. In the Service Instances table, select the service instance for which you want to upgrade the service design version.

3. Click **More > Upgrade Service Design**.

4. On the Upgrade Service Design dialog box, click the drop-down to select the service design version you want to upgrade to.

5. Click **Yes** to confirm.

You can view the status of the service design upgrade in the State column of the Service Instances table. The state is displayed as **Processing** when the upgrade is in process. After successful upgrade, the state is displayed as **Success**.

SEE ALSO

[About the Service Designs Page](#) | 653

Upload a Customized Service Design

You use a service design to add and provision a service in the network. Routing Director allows you to upload customized service designs and use them to provision corresponding services in the network.



NOTE: Uploading a customized service design is a beta feature in this release.

To create and customize a service design, you must contact [Juniper Networks Professional Services](#). The Juniper Networks Professional Services team provides you with a customized service design file in the .tgz format. Additionally, they will also provide you with common.tgz and network-resource.tgz files with corresponding configurations that map with the customized service design.

To upload the customized service design to the GUI (**Orchestration > Service > Service Catalog**), perform the following steps:

1. Log in to the Paragon Shell CLI as the root user.
2. Type exit to exit Paragon Shell to the Linux root shell.
3. Start a cMGD pod using the following command:

```
root@node:~# kubectl -n foghorn scale --replicas=1 deployments/cmgd
deployment.apps/cmgd scaled
```

4. Launch the service orchestration cMGD CLI bash shell using the following command:

```
root@node:~# kubectl -n foghorn exec -it -c cmgd deployment/cmgd -- bash
root@cmgd-6866d5d9c9-4bwsr>
```

You are now logged in to the service orchestration cMGD CLI bash shell.

5. Copy the customized service design .tgz, common.tgz, and network-resource.tgz files to the service orchestration cMGD CLI. For instance in this example, we have copied the customized service design CSD.tgz, common.tgz and network-resource.tgz files to the /data-models/projects folder. Here, CSD.tgz is the customized service design file and common.tgz and network-resource.tgz are the dependency files.

6. Type `cli` and press **Enter** to launch the CLI mode.
7. View organizations available in Routing Director to add the customized service design.

```
root@cmgd-6866d5d9c9-4bwsr> show paragon orgs
Org-Name          Org-ID
SD-test           cd71dd08-72bb-4697-b516-655de87f69c6
```

Where:

Org-Name is the name of the organization created in Routing Director.

Org-ID is the unique auto-generated alphanumeric ID assigned to the organization.

8. Enter the configuration mode.

```
root@cmgd-6866d5d9c9-4bwsr> configure
Entering configuration mode

[edit]
```

9. Configure the organization ID to set the context of all tasks required to add the customized service design.

```
root@cmgd-6866d5d9c9-4bwsr# set foghorn:core org-id cd71dd08-72bb-4697-b516-655de87f69c6

[edit]
```

Where:

cd71dd08-72bb-4697-b516-655de87f69c6 is the unique Org-ID.

10. Commit the changes and exit the configuration mode.

```
root@cmgd-6866d5d9c9-4bwsr# commit and-quit
commit complete
Exiting configuration mode
```

11. Add projects to the service catalog for all organizations. Projects contain a service design and its dependency files required to provision a service.

- Add common files to the service catalog. Common files are reusable components required by multiple service designs.

```
root@cmgd-6866d5d9c9-4bwsr> request service catalog project add /data-models/projects/
common.tgz
```

Where `common.tgz` file contains common files that includes standard YANG models, Jinja templates, third-party schemas, shared JQS, party files, standard IETF templates, YANG models, and so on.

- Add network resource configuration files to the service catalog. The network resource file contains network resource service designs that acts as producers of shared resources and are consumed by service designs.

```
root@cmgd-6866d5d9c9-4bwsr> request service catalog project add /data-models/projects/
network-resource.tgz
```

Where `network-resource.tgz` file contains definitions of network elements such as IP address pools, VLAN IDs, VPN route distinguishers, and so on.

- Add your customized service design `CSD.tgz`, `common.tgz`, and `network-resource.tgz` to the service catalog.

```
root@cmgd-6866d5d9c9-4bwsr> request service catalog project add /data-models/projects/
CSD.tgz
```

```
root@cmgd-6866d5d9c9-4bwsr> request service catalog project add /data-models/projects/
common.tgz
```

```
root@cmgd-6866d5d9c9-4bwsr> request service catalog project add /data-models/projects/
network-resource.tgz
```

12. Verify the projects that you added to the service catalog using the following command:

```
root@cmgd-6866d5d9c9-4bwsr> show service catalog projects
```

Project Name	Project ID
common	1670dc07-c891-4031-a15c-2885939f75c7

```

DesignID: iana Version: 0.1.1
DesignID: ietf Version: 0.1.1
DesignID: ietf-draft Version: 0.1.1
DesignID: juniper Version: 0.1.1

network-resource          d090c91c-3163-4b3d-ba43-cba9b824a0e9
  DesignID: network-design Version: 0.1.9
  DesignID: l2-addr Version: 0.1.5
  DesignID: l3-addr Version: 0.2.6
  DesignID: routing Version: 0.1.9
  DesignID: topo Version: 0.6.25
  DesignID: vpn-resources Version: 0.3.5

CSD                        e5c6a17f-7b24-4136-908a-145e80b9d898
  DesignID: csd Version: 0.0.1

<Output Snipped>

```

Where:

Project Name displays the names of projects that you added to the service catalog.

Project ID displays the unique ID assigned to different projects available in the service catalog.

13. (Optional) Log in to the Routing Director GUI and navigate to the Service Designs page (**Orchestration > Service Catalog**). The newly added service design version 0.0.1 must be listed in the **Pending Version** column.
14. Install all the files that you added to the service catalog using the following command:

```

root@cmgd-6866d5d9c9-4bwsr> request service project add common_1670dc07-c891-4031-
a15c-2885939f75c7

```

```

root@cmgd-6866d5d9c9-4bwsr> request service project add network-resource_d090c91c-3163-4b3d-
ba43-cba9b824a0e9

```

```

root@cmgd-6866d5d9c9-4bwsr> request service project add
CSD_e5c6a17f-7b24-4136-908a-145e80b9d898

```

15. This command takes sometime to complete.

A warning message confirms a successful update. Restart the CLI to complete installation of the new version. Confirm **Yes**, when prompted.

```
root@cmgd-6866d5d9c9-4bwsr> request service project add
CSD_e5c6a17f-7b24-4136-908a-145e80b9d898

root@cmgd-6866d5d9c9-4bwsr>

WARNING: cli has been replaced by an updated version:
CLI release 20240313.184005_ningarajs.r1408131 built by ningarajs on 2024-03-13 19:02:49 UTC
Restart cli using the new version ? [yes,no] (yes) yes

Restarting cli ...
root@cmgd-6866d5d9c9-4bwsr>
```

16. Now the customized service design is installed and must be available in the GUI. Log in again to the Routing Director GUI and the newly installed service design version 0.0.1 must be listed as **Default** under the **Version** column on the Service Designs page (**Orchestration > Service Catalog**). For more information, see ["About the Service Designs Page" on page 653](#).

You can create and provision a service based on the service design that you added. The **Add** drop-down on the Service Instances page (**Orchestration > Service Instances**) displays the name of the service design that you installed. Select the service design to create a corresponding service instance. For more information, see ["About the Service Instances Page" on page 829](#).

Manage Customers

IN THIS CHAPTER

- [About the Customer Inventory Page | 664](#)
- [Add a Customer | 666](#)
- [Edit and Delete Customers | 667](#)

About the Customer Inventory Page

IN THIS SECTION

- [Tasks You Can Perform | 664](#)
- [Field Descriptions | 665](#)

Routing Director allows a service provider to create and manage an inventory of their customers. Customer here refers to a user or an organization that utilizes a service from a service provider. The Customer Inventory page lists customers by name, customer ID, the service instances provisioned for the customer, customer reference number, and description.

After you enter your customer details in the customer inventory, you can reference the customer name from the inventory when you create a service instance.

To access the Customer Inventory page, click **Orchestration** > **Customers** on the navigation menu.

Tasks You Can Perform

You can perform the following tasks on the Customer Inventory page:

- View customer details—The customer table displays the name, reference number, and a description of the customer; see [Table 156 on page 665](#) for details.

- Add a new customer to the customer inventory; see ["Add a Customer" on page 666](#).
- Modify or delete customer details from the customer inventory; see ["Edit and Delete Customers" on page 667](#).

-  **NOTE:** You cannot delete a customer if you have services provisioned for the customer. You must deprovision services before deleting a customer associated with the service.
- View services provisioned for a customer—To see the services instances you have provisioned for a customer, select the customer and click **View Instances** under the Service Instances header. You are directed to the Service Instances page that displays the service instances that you provisioned for the customer.
- Filter the data displayed in the table—Click the filter icon (funnel) and select whether you want to show or hide advanced filters. You can then add or remove filter criteria, save criteria as a filter, apply or clear filters, and so on. The filtered results are displayed on the same page.
- Search by using keywords—Click the search icon (magnifying glass), enter the search term in the text box, and press Enter. The search results are displayed on the same page.
- Sort, resize, or re-arrange columns in a table (grid).



NOTE: Sort functionality is available only for the Names column.

Field Descriptions

[Table 156 on page 665](#) lists the fields on the Customer Inventory page.

Table 156: Fields on the Customer Inventory Page

Field	Description
Name	Name of the customer.
Customer ID	Unique alphanumeric ID automatically assigned to a customer after you enter your customer details in the customer inventory.
Service Instances	Click the View Instances hyperlink to view the service instances provisioned for the customer.

Table 156: Fields on the Customer Inventory Page (Continued)

Field	Description
Reference Number	Unique number that you (service provider) assign a customer. This number might be internally referenced from the Customer Relationship Management (CRM) system.
Description	Description about the customer.

Add a Customer

A customer refers to a user or an organization that utilizes services of a service provider. A Super User or a Network Admin can add customers to Routing Director for provisioning services such as L3VPN.

To add a customer:

1. Click **Orchestration > Customers**.
The Customer Inventory page appears.
2. Click the add (+) icon.
The Add Customers page appears.
3. Enter values by referring to [Table 157 on page 666](#).
4. Click **OK**.

A message indicating that the customer is created appears. You are returned to the Customer Inventory page where you can view the newly added customer listed in the Customer Inventory table.

Table 157: Fields on the Add Customer Page

Field	Description
Name	Enter a name for the customer. The customer name must be unique within an organization.

Table 157: Fields on the Add Customer Page *(Continued)*

Field	Description
Reference Number	Enter a reference number for the customer. This number might be internally referenced by the Customer Relationship Management (CRM) system.
Description	Enter a description for the customer.

RELATED DOCUMENTATION

- [Edit and Delete Customers | 667](#)
- [About the Customer Inventory Page | 664](#)

Edit and Delete Customers

IN THIS SECTION

- [Edit Customer Details | 667](#)
- [Delete a Customer | 668](#)

Use the procedures in this topic to edit and delete customer details from the customer inventory.

Edit Customer Details

To edit customer details:

1. Click **Orchestration > Customers**.
The Customer Inventory page appears.
2. Click the **Edit** (Pen) icon.
The Edit Customer page appears.
3. Enter values by referring to [Table 157 on page 666](#).

4. Click **OK**.

A message indicating that the customer detail is updated appears. You are returned to the Customer Inventory page where you can view the updates made to the customer details.

Delete a Customer



NOTE: You cannot delete a customer if you have services provisioned for the customer. You must remove the provisioned services before deleting a customer.

To delete a customer:

1. Click **Orchestration > Customers**.

The Customer Inventory page appears.

2. Click the **Delete** (Dustbin) icon.

A Confirm Delete page appears.

3. Click **Yes** to confirm deleting the customer.

A confirmation message appears informing that the customer is deleted and removed from the Customer Inventory table.

SEE ALSO

[Add a Customer | 666](#)

[About the Customer Inventory Page | 664](#)

Add Resources for Network Services

IN THIS CHAPTER

- [Resource Instance Overview | 669](#)
- [Add Network Resources for Provisioning Services | 670](#)
- [About the Resource Instances Page | 672](#)
- [Add a Resource Instance | 674](#)
- [Add Network Resource Pools | 676](#)
- [Configure Topology Resources in Network Implementation Plan | 743](#)
- [Configure Resource Pools for Resource Instances | 745](#)
- [Modify, Delete, and Restore Resource Instances | 766](#)
- [Manage Placement Configurations for Service Instances | 768](#)
- [View Network Resources | 771](#)

Resource Instance Overview

To provision a service in the network, you (superuser or network administrator) must configure the network resource pools for the service. A resource instance defines the elements of a network resource pool. For example, a VPN resource instance defines resources such as route distinguishers, route targets, and so on for provisioning L3VPN, EVPN, and L2 circuit services. You can create the following types of resource instances based on the guidelines and templates defined by the corresponding resource designs that Routing Director provides:

- Routing resource instance based on the **routing** resource design—Defines the routing resource pool (AS, route reflector, and so on) required to provision the infrastructure service.
- Layer 3 resource instance based on the **I3-addr** resource design—Defines the Layer 3 resource pool (IPv4 and loopback addresses) to provision L3VPN and infrastructure services.
- Layer 2 resource instance based on the **I2-addr** resource design—Defines the Layer 2 resource pool (EVPN ESIs, LACP administrative keys and system IDs) to provision EVPN services.

- Topology resource instance based on the **topo** resource design—Defines the topology resource pool (PE details, PE to CE connection details, VLANs, and so on) to provision L3VPN, EVPN, and EVPN-VPWS services.

Since you can also manage resources from within network implementation plans, a blank topology **topo** instance is auto created on the Resources Instances page (**Orchestration > Service > Resource Instances**). This topology instance is a read-only instance and is owned by the network operator.

- VPN resource instance based on the **vpn-resources** resource design—Defines the VPN resource pool (route distinguishers, route targets, and so on) to provision L3VPN, EVPN, and L2 circuit services.

You can also modify and delete a resource instance. When you create, modify, or delete a resource instance, and then save and commit it, you generate a corresponding service order (create, modify, or delete) that activates the automated workflows associated with the service order. You can monitor the execution status of a service order. After a create or modify workflow is successfully executed, the resource pool is uploaded to the Routing Director database. When you provision a service, you use the uploaded network resource configurations to assign and allocate network resources for the service. The delete workflow deletes or removes the resource pools from the database.

The Resource Instances page (**Orchestration > Service > Resource Instances**) lists all the resource instances created in an organization.

RELATED DOCUMENTATION

[About the Resource Instances Page | 672](#)

[Add a Resource Instance | 674](#)

[Configure Resource Pools for Resource Instances | 745](#)

[Modify, Delete, and Restore Resource Instances | 766](#)

Add Network Resources for Provisioning Services

As a prerequisite to provisioning a network service, you (superuser or network administrator) must configure and upload network resource pools for the service. Routing Director provides the following network resource designs to configure network resource pools for provisioning network services:

- Routing and L3 resource designs to configure network resource pools for provisioning infrastructure services.
- L2 resource design to configure network resource pools defining EVPN ESI details, LACP administrative keys and system identifiers for provisioning EVPN services.

- Topology resource design to configure resource pools defining site details for provisioning L3VPN, EVPN, and EVPN-VPWS services.
- VPN-resources resource design to configure resource pools defining route distinguishers, route targets, and virtual circuit IDs for provisioning L3VPN, EVPN, and L2 circuit services.

You can create, modify, or delete resource pools for service provisioning. Each of these operations (create, modify, delete) generates a service order to activate the corresponding workflow. The create and modify workflows upload new or modified resource pools to the Routing Director database. The delete workflow deletes or removes a resource pool from the database.

You can upload resource pools by using any of the following options:

- Upload network resources by using the Resource Instances page (**Orchestration > Service > Resource Instances**) of the Routing Director GUI. See ["About the Resource Instances Page" on page 672](#).
- Upload network resources by using REST APIs. See ["Add Network Resource Pools by Using REST APIs" on page 677](#).
- Upload network resources by using the service orchestration cMGD CLI.

To add network resources by using the service orchestration cMGD CLI:

1. Log in to the service orchestration cMGD CLI. See ["Access the Service Orchestration cMGD CLI" on page 1577](#).
 2. Execute the request service order load *<service-order-name>* command to upload a new or modified network resource service order to the service orchestration cMGD environment. See ["request service order load" on page 1633](#).
 3. Execute the request network resources load ? command to view all the network resources service orders available in service orchestration cMGD.
 4. Execute the request network resources load *<service-order-name>* command to upload the network configurations you created or modified in the service order to the Routing Director database. See ["request network resources load" on page 1631](#).
- Upload network resources by configuring the resources in a network implementation plan.

When you use a network implementation plan to onboard a device, the resources defined in the network implementation plan are reserved on the device. You can provision a service by using the reserved resource immediately after the device is onboarded.

See ["Configure Topology Resources in Network Implementation Plan" on page 743](#).

About the Resource Instances Page

IN THIS SECTION

- [Tasks You Can Perform | 672](#)
- [Field Descriptions | 673](#)

Before provisioning a service in the network using Routing Director, you must configure and upload the network resource pools for the service. The Resource Instances page lists the resource instances that you (service provider) create, modify, or delete to provision a service for your customers. The resource instances are listed by name, the associated resource design, and so on.

To access the **Resource Instances** page, click **Orchestration > Resource Instances**.

Tasks You Can Perform

You can perform the following tasks on this page:

- View details of resource instances.

Use one of the following options to view resource instance details on the right pane:

- Hover over the resource instance and click the **Details** icon that appears next to the resource instance name.
- Right-click the resource instance name and click **Show Detail**.
- Select the resource instance and click **More > Show Detail**.

The resource instance details pane appears on the right side of the page. For information about the fields on the pane, see [Table 158 on page 673](#). To view the resource instance details in JSON format, click the **JSON View** hyperlink on the details pane.

- Add a new resource instance. See ["Add a Resource Instance" on page 674](#).
- Modify a resource instance. See ["Modify a Resource Instance" on page 766](#).
- Delete or deprovision a resource instance. See ["Delete a Resource Instance" on page 767](#).
- Search by using keywords—Click the search icon (magnifying glass), enter the search term in the text box, and press Enter. The search results are displayed on the same page.

- Show or hide columns in the table or reset page preferences, using the vertical ellipsis menu.
- Sort, resize, or re-arrange columns in a table (grid).
- Set the number of rows to be displayed in the table—Click the **Display** drop-down at the bottom-right of the Resource Instances table to set the number of rows (15, 30, 50, or 100) to be displayed in the table. The default number of rows displayed is 15.

Use the arrow heads to traverse through the pages or enter a page number to view the entries listed on that page.

- Refresh data displayed in the Resource Instances table—Click the Refresh icon at the bottom-left corner of the Resource Instances table to view the most recent updates.

Field Descriptions

Table 158 on page 673 lists the fields on the Resource Instance page.

Table 158: Fields on the Resource Instance Page

Field	Description
Name	Name of the resource instance. For example, topology. You can create resource instances only for the network-operator customer using the GUI. When the customer is not network-operator, the name of the customer is displayed separated by a colon from the resource instance name in the resource_name : customer_name format. For example, topology : test-abc, where test-abc is the customer name. To create resource instances for customers other than network-operator, use APIs.
Resource Design	Name of the resource design associated with the resource instance. For example, topology (Topo) resource design associated with topology resource instances for L3VPN, EVPN, EVPN-VPWS, or L2 circuit services.
Design Version	Version number of the resource design associated with the resource instance. For example, 0.2.4.
Status Message	Status of the resource instance.

Table 158: Fields on the Resource Instance Page *(Continued)*

Field	Description
Workflow Run ID	ID of the workflow run associated with the resource instance.
Last Modified	The date and time when the resource instance was last modified.
Last Modified By	Name of the user who last modified the resource instance.

RELATED DOCUMENTATION

[Add a Resource Instance | 674](#)

[Configure Resource Pools for Resource Instances | 745](#)

[Modify, Delete, and Restore Resource Instances | 766](#)

Add a Resource Instance

You can create a resource instance based on the corresponding resource design and configure the resources available in the resource pool to provision a service. After you add a resource instance and configure the resource pool for the instance, you can save and commit the instance. When you save and commit a resource instance, Routing Director activates a workflow run to upload the resource instance to the database. You can monitor the workflow execution state by viewing the workflow run details.

A resource instance is created for a customer to define the available network resources and to allocate these resources when provisioning a VPN service for the customer. You can add resources instances only for the default **network-operator** customer from the GUI.

To add a resource instance:

1. Click **Orchestration > Resource Instances**.
The Resource Instances page appears.
2. Click the add (+) icon above the Resource Instances table.
The Add New Resource Instance page appears.
3. Enter the values for instance name and select a resource design by referring to [Table 159 on page 675](#).
4. Do one of the following:

- Click **Create**.

A template of the resource instance based on the corresponding resource design is created. The instance name is listed in the Resource Instances table. You are directed to the Modify *Resource-Instance-Name* page, which displays the following panes:

- The Resource Editor on the left pane.
- The resource Inventory in JSON format on the right pane.

The inventory lists all network resources associated with the current resource instance, that are available in the Routing Director database. You can expand or collapse the sections to view resource configuration details.

- Click **Close** to discard your changes and exit the page. To restart the process, go to [1](#).
5. Configure the network resource pools as described in ["Configure Resource Pools for Resource Instances" on page 745](#).
 6. Click **Proceed** to save the resource instance configuration.
The Compare Resource Definition page appears.
 7. Click **Save and Commit**.
A service order is created and the Monitor order status page appears and displays the execution status of the service order. The resource instance is first saved and then committed.
 8. (Optional) If the service order commit fails, navigate to the *workflow-ID* page to debug and fix the issue.
 9. Click **OK**.
The Resource Instances page appears and displays the newly added resource instance.

Table 159: Fields on the Add New Resource Instance Page

Field	Description
Instance Name	The name of the resource instance. For example, vpn1.

Table 159: Fields on the Add New Resource Instance Page (*Continued*)

Field	Description
Resource Design	Routing Director provides the following resource designs: • Topology—The topology resource design provides guidelines to create topology resource pools for L3VPN, EVPN, EVPN-VPWS, and L2 circuit services. • VPN Resources—The VPN Resources resource design provides guidelines to create VPN resource pools for L3VPN, EVPN, and L2 circuit services. • Routing—The routing resource design provides guidelines to create routing resource pools for infrastructure services. • L3-Addr—The L3 resource design provides guidelines to create L3 resource pools for infrastructure services. • L2-Addr—The L2 resource design provides guidelines to create L2 resource pools for EVPN services.

RELATED DOCUMENTATION

[Resource Instance Overview](#) | 669

[Configure Resource Pools for Resource Instances](#) | 745

[Modify, Delete, and Restore Resource Instances](#) | 766

Add Network Resource Pools

IN THIS SECTION

- [Upload Network Resource Pools for Device Onboarding in JSON Format](#) | 677
- [Add Network Resource Pools by Using REST APIs](#) | 677
- [Sample Files](#) | 680

A network resource pool defines values for the network resources, such as IPv4 loopback addresses, interface IP addresses, segment identifiers (SIDs), BGP cluster IDs, and so on, that are assigned to devices in your network.

Routing Director assigns values to the network resources in a device profile and an interface profile when automatic configuration is enabled for the network resources.

You can create a resource pool by:

- Uploading a JSON file. See ["Upload Network Resource Pools for Device Onboarding in JSON Format" on page 677](#).
- Manually entering the values. See ["Add a Resource Instance" on page 674](#).
- Using REST APIs. See ["Add Network Resource Pools by Using REST APIs" on page 677](#).

Upload Network Resource Pools for Device Onboarding in JSON Format

To upload network resource pools for device onboarding in JSON format:

1. Navigate to **Orchestration > Service > Resource Instances**.
The Resource Instances page appears.
2. Click **More > Download Sample Network Resources** to download sample files that define the resource pools.
Sample network resource files **l3-addr.json**, **topo_sample.json**, and **routing.json** files are downloaded to your local system.

The **l3-addr.json** file defines the resource pools for loopback address and IPv4 addresses. The **routing.json** file defines the resource pools for autonomous system (AS) number, SIDs, and BGP cluster IDs.
3. Define your network resource files by editing the values for network resources in the sample files.
4. Save your network resource files.
5. Click **More > Upload Network Resources** to upload the files and create the resource pools.
6. In the browser dialog box, browse for the network resource files and click **Upload** to upload the files.
7. (Optional) Click **More > View Network Resources** to view the resource pools that you uploaded and that are available in Routing Director. See ["View Network Resources" on page 771](#).

After you define the resource pools, you can add device profiles and interface profiles to Routing Director. See ["Add a Device Profile" on page 208](#) and ["Add an Interface Profile" on page 222](#) for details.

Add Network Resource Pools by Using REST APIs

To create network resource pools by using REST APIs, you should be familiar working with the tools such as Postman to make API requests to Routing Director. To view all REST API documentation and

OpenAPI definitions, navigate to **Help (?) > API Docs** in the GUI. All service orchestration APIs are defined in the **Service Orchestration** section.

As resource instances are a specific type of service instance, the relevant APIs to create, modify and delete resource instances are the following Order APIs:

- Upload Service Order
- Execute Service Instance Workflow
- Get Service Instance
- Get Order History for SI

To see the resulting resources, as used during placement, the placement API **Get Network Resources** must be used.

Example use of the API, including payload using Postman is described in this topic.

You will need values for the following parameters to create resource pools for the network resources:

- The URL to the environment where Routing Director is running.
- ID of the organization where you want to add the resource pools.
- Username for accessing the organization.
- Password for accessing the organization.

To create a resource pool:



NOTE: This procedure details how to use Postman to execute the REST APIs. You can use other tools as well to execute the REST APIs.

1. Download the Postman application from <https://www.postman.com/downloads/>.
2. Install and configure Postman on your system.
For information about working with Postman application, see [Postman documentation](#).
3. Create a Postman environment file.
For information about creating an environment file, see <https://learning.postman.com/docs/sending-requests/managing-environments/>. See "Sample Postman Environment File" on page 680 for a sample of the Postman environment file.
4. Create a Postman collection file.
 - For device onboarding, see "Sample Postman Collection File for Device Onboarding" on page 681 for a sample of the Postman collection file and Table 160 on page 692 for the REST APIs

included in the sample Postman collection file. The sample collection file includes APIs for creating resource pools for IPv4 addresses and BGP cluster IDs.

- For service orchestration, see ["Sample Postman Collection Files for Service Orchestration" on page 695](#) for a sample of the Postman collection file and [Table 161 on page 702](#) for the REST APIs included in the sample Postman collection file. The sample collection file includes APIs for creating the VPN and topology resource pools for the L3VPN service.
5. Execute the REST API for getting credentials to access Routing Director and get organization ID. You need the organization ID to update the *organization_id* parameter in the environment file (*ORG* in the ["Sample Postman Environment File" on page 680](#)).

In the ["Sample Postman Collection File for Device Onboarding" on page 681](#) and ["Sample Postman Collection Files for Service Orchestration" on page 695](#), the REST API to be executed for getting the organization ID is 01-who am i and get orgs. A snippet of the sample response for the 01-who am i and get orgs API request is as follows:

```
{ "scope": "org",

"org_id": "3b1c3556-5c05-4abc-9bf4-3bdd9c231f23",
"role": "admin",
"name": "TestingTasks" }
```

Alternatively, access the Organization Settings page (**Settings menu > System Settings** on the banner) on the Routing Director GUI to get the organization ID from the Organization ID field.

6. In the environment file, ensure that:
 - Variables *User* and *Password* are set to your username and password used for logging in to the Routing Director GUI.
 - Verify that the *server* is set to the Routing Director GUI VIP address `web-ui-vip-address`.
7. Import the environment file into Postman.
8. Import the collection file into Postman.
9. Execute the REST APIs in the collection file to create resource pools.
10. After the APIs complete execution and return a response indicating that the resource pools are created, view the network resources added to Routing Director.
 - For device onboarding, see ["View Network Resources" on page 771](#).
 - For service orchestration:
 - a. Click **Orchestration > Instances**.

The Service Instances page appears.

- b. Click **More > View Network Resources** on the Service Instances page to view the resource pools that you uploaded and that are available in Routing Director.

Sample Files

IN THIS SECTION

- [Sample Postman Environment File | 680](#)
- [Sample Postman Collection File for Device Onboarding | 681](#)
- [Sample REST API to Create an IPv4 Address Pool | 692](#)
- [Sample REST API to Create BGP Cluster ID Pool | 693](#)
- [Sample Postman Collection Files for Service Orchestration | 695](#)
- [Sample REST API to Create Topology Resources for L3VPN and L2VPN Services | 739](#)
- [Sample REST API to Create VPN Resources for L3VPN, L2VPN, and L2 Circuit Services | 742](#)

This section provides a sample of the environment file, collection file, and the list of REST APIs that you can use to define network resources pools.

Sample Postman Environment File

The following is a sample Postman environment file.

```
{
  "id": "dae981a2-da91-4d6f-9094-87e6ea05003c",
  "name": "00-00-pa",
  "values": [
    {
      "key": "server",
      "value": "web-ui-vip-address",
      "enabled": true
    },
    {
      "key": "port",
      "value": "443",
      "enabled": true
    }
  ],
}
```

```

    {
      "key": "Password",
      "value": "abc123",
      "type": "secret",
      "enabled": true
    },
    {
      "key": "User",
      "value": "user@abc.com",
      "enabled": true
    },
    {
      "key": "ORG",
      "value": "34a55586-2baf-4cce-b2e0-0b293b223af1",
      "type": "default",
      "enabled": true
    },
    {
      "key": "SITE_ID",
      "value": "",
      "type": "any",
      "enabled": true
    }
  ],
  "_postman_variable_scope": "environment",
  "_postman_exported_at": "2023-04-20T12:04:35.537Z",
  "_postman_exported_using": "Postman/10.12.13"

```

Sample Postman Collection File for Device Onboarding

The following is a sample Postman collection file to define values for IPv4 addresses and BGP cluster IDs.

```

{
  "info": {
    "_postman_id": "7a32e6b1-d4ca-4166-9a9f-3777c2ae6ce4",
    "name": "Resource Profile creation",
    "schema": "https://schema.getpostman.com/json/collection/v2.1.0/collection.json",
    "_exporter_id": "829664"
  },
  "item": [

```

```

{
  "name": "01-Who am I and get orgs",
  "event": [
    {
      "listen": "prerequisite",
      "script": {
        "exec": [
          "var x=CryptoJS.enc.Utf8.parse(postman.getEnvironmentVariable(\"User\"))+\"\":\"+postman.getEnvironmentVariable(\"Password\");",
          "var authHeader=CryptoJS.enc.Base64.stringify(x);",
          "pm.request.headers.add({key: \"Authorization\", value: \"Basic\""+authHeader});",
          ""
        ],
        "type": "text/javascript"
      }
    },
    {
      "listen": "test",
      "script": {
        "exec": [
          "var jsonData = JSON.parse(responseBody);",
          "postman.setEnvironmentVariable(\"ORG\", jsonData.privileges[0].org_id);",
          ""
        ],
        "type": "text/javascript"
      }
    }
  ],
  "request": {
    "method": "GET",
    "header": [],
    "url": {
      "raw": "https://{{server}}:{{port}}/api/v1/self",
      "protocol": "https",
      "host": [
        "{{server}}"
      ],
      "port": "{{port}}",
      "path": [
        "api",
        "v1",
        "self"
      ]
    }
  }
}

```

```

        ]
      }
    },
    "response": []
  },
  {
    "name": "02-pick-site",
    "event": [
      {
        "listen": "prerequest",
        "script": {
          "exec": [
            "var x=CryptoJS.enc.Utf8.parse(postman.getEnvironmentVariable(\"User\"))+\"\\\":\"+postman.getEnvironmentVariable(\"Password\\\")\";",
            "var authHeader=CryptoJS.enc.Base64.stringify(x);",
            "pm.request.headers.add({key: \"Authorization\", value: \"Basic\""+authHeader});",
            ""
          ],
          "type": "text/javascript"
        }
      },
      {
        "listen": "test",
        "script": {
          "exec": [
            "var jsonData = JSON.parse(responseBody);",
            "postman.setEnvironmentVariable(\"SITE_ID\", jsonData[0].id);"
          ],
          "type": "text/javascript"
        }
      }
    ],
    "request": {
      "method": "GET",
      "header": [],
      "url": {
        "raw": "https://{{server}}:{{port}}/api/v1/orgs/{{ORG}}/sites",
        "protocol": "https",
        "host": [
          "{{server}}"
        ],
        "port": "{{port}}",

```

```

        "path": [
            "api",
            "v1",
            "orgs",
            "{{ORG}}",
            "sites"
        ]
    },
    "response": []
},
{
    "name": "03-Create L3 Addr",
    "event": [
        {
            "listen": "prerequisite",
            "script": {
                "exec": [
                    "var x=CryptoJS.enc.Utf8.parse(postman.getEnvironmentVariable(\"User\")+\":\")+postman.getEnvironmentVariable(\"Password\");",
                    "var authHeader=CryptoJS.enc.Base64.stringify(x);",
                    "pm.request.headers.add({key: \"Authorization\", value: \"Basic\","+authHeader});",
                    ""
                ],
                "type": "text/javascript"
            }
        },
        {
            "listen": "test",
            "script": {
                "exec": [
                    ""
                ],
                "type": "text/javascript"
            }
        }
    ],
    "request": {
        "method": "POST",
        "header": [],
        "body": {
            "mode": "raw",

```

```

        "raw": "{\n    \"customer_id\": \"network-operator\",\n    \"design_id\":  
    \"l3-addr\",\n    \"instance_id\": \"l3-stuff\",\n    \"operation\": \"create\",\n    \"l3_addr  
    \": {\n        \"loopbacks\": [\n            {\n                \"name\":  
                \"range-192\",\n                \"prefix\": \"10.1.192.0/18\"\n            },\n            {\n                \"name\": \"range-2\",\n                \"prefix\":  
                \"10.2.2.0/24\"\n            },\n            {\n                \"name\":  
                \"range-3\",\n                \"prefix\": \"10.3.3.0/24\"\n            }\n        ],\n        \"ipv4_prefixes\": [\n            {\n                \"name\": \"pool-11\",\n                \"prefix\": \"10.11.11.0/24\"\n            },\n            {\n                \"name\":  
                \"pool-12\",\n                \"prefix\": \"10.12.12.0/24\"\n            },\n            {\n                \"name\": \"pool-13\",\n                \"prefix\":  
                \"10.13.13.0/24\"\n            },\n            {\n                \"name\":  
                \"pool-14\",\n                \"prefix\": \"10.14.14.0/24\"\n            },\n            {\n                \"name\": \"pool-15\",\n                \"prefix\":  
                \"10.15.15.0/24\"\n            },\n            {\n                \"name\":  
                \"pool-16\",\n                \"prefix\": \"10.16.16.0/24\"\n            },\n            {\n                \"name\": \"pool-17\",\n                \"prefix\":  
                \"10.17.17.0/24\"\n            },\n            {\n                \"name\":  
                \"pool-18\",\n                \"prefix\": \"10.18.18.0/24\"\n            },\n            {\n                \"name\": \"pool-19\",\n                \"prefix\":  
                \"10.19.19.0/24\"\n            },\n            {\n                \"name\":  
                \"pool-20\",\n                \"prefix\": \"10.20.20.0/24\"\n            },\n            {\n                \"name\":  
                \"pool-21\",\n                \"prefix\": \"10.21.21.0/24\"\n            },\n            {\n                \"name\": \"pool-22\",\n                \"prefix\":  
                \"10.22.22.0/24\"\n            },\n            {\n                \"name\":  
                \"pool-23\",\n                \"prefix\": \"10.23.23.0/24\"\n            },\n            {\n                \"name\": \"pool-24\",\n                \"prefix\":  
                \"10.24.24.0/24\"\n            },\n            {\n                \"name\":  
                \"pool-25\",\n                \"prefix\": \"10.25.25.0/24\"\n            },\n            {\n                \"name\": \"pool-26\",\n                \"prefix\":  
                \"10.26.26.0/24\"\n            },\n            {\n                \"name\":  
                \"pool-27\",\n                \"prefix\": \"10.27.27.0/24\"\n            },\n            {\n                \"name\": \"pool-28\",\n                \"prefix\":  
                \"10.28.28.0/24\"\n            },\n            {\n                \"name\":  
                \"pool-29\",\n                \"prefix\": \"10.29.29.0/24\"\n            },\n            {\n                \"name\": \"pool-30\",\n                \"prefix\":  
                \"10.30.30.0/24\"\n            },\n            {\n                \"name\":  
                \"pool-31\",\n                \"prefix\": \"10.31.31.0/24\"\n            },\n            {\n                \"name\":  
                \"pool-32\",\n                \"prefix\": \"10.32.32.0/24\"\n            },\n            {\n                \"name\": \"pool-33\",\n                \"prefix\":  
                \"10.33.33.0/24\"\n            }\n        ]\n    },\n    \"options\": {\n        \"raw\": {

```

```

        "language": "json"
      }
    },
    "url": {
      "raw": "https://{{server}}:{{port}}/service-orchestration/api/v1/orgs/
{{ORG}}/order",
      "protocol": "https",
      "host": [
        "{{server}}"
      ],
      "port": "{{port}}",
      "path": [
        "service-orchestration",
        "api",
        "v1",
        "orgs",
        "{{ORG}}",
        "order"
      ]
    },
    "response": []
  },
  {
    "name": "04-Exec L3 Addr",
    "event": [
      {
        "listen": "prerequisite",
        "script": {
          "exec": [
            "var x=CryptoJS.enc.Utf8.parse(postman.getEnvironmentVariable(\"User
\")+\":\")+postman.getEnvironmentVariable(\"Password\"));",
            "var authHeader=CryptoJS.enc.Base64.stringify(x);",
            "pm.request.headers.add({key: \"Authorization\", value: \"Basic
\"+authHeader});",
            ""
          ],
          "type": "text/javascript"
        }
      },
      {
        "listen": "test",

```

```

        "script": {
            "exec": [
                ""
            ],
            "type": "text/javascript"
        }
    },
    ],
    "request": {
        "method": "POST",
        "header": [],
        "url": {
            "raw": "https://{{server}}:{{port}}/service-orchestration/api/v1/orgs/
{{ORG}}/order/customers/network-operator/instances/l3-addr/exec",
            "protocol": "https",
            "host": [
                "{{server}}"
            ],
            "port": "{{port}}",
            "path": [
                "service-orchestration",
                "api",
                "v1",
                "orgs",
                "{{ORG}}",
                "order",
                "customers",
                "network-operator",
                "instances",
                "l3-addr",
                "exec"
            ]
        }
    },
    "response": []
},
{
    "name": "05-Create Routing Resources",
    "event": [
        {
            "listen": "prerequisite",
            "script": {
                "exec": [

```



```

        "var x=CryptoJS.enc.Utf8.parse(postman.getEnvironmentVariable(\"User
\")+\":\")+postman.getEnvironmentVariable(\"Password\");",
        "var authHeader=CryptoJS.enc.Base64.stringify(x);",
        "pm.request.headers.add({key: \"Authorization\", value: \"Basic
\"+authHeader});",
        ""
    ],
    "type": "text/javascript"
}
},
{
    "listen": "test",
    "script": {
        "exec": [
            ""
        ],
        "type": "text/javascript"
    }
}
],
"request": {
    "method": "POST",
    "header": [],
    "body": {
        "mode": "raw",
        "raw": "{\n  \"customer_id\": \"network-operator\",\n  \"design_id\":
\"routing\",\n  \"instance_id\": \"routing-stuff\",\n  \"operation\": \"create\",\n
\"routing\": {\n    \"autonomous_system\": [\n      {\n        \"name\":
65200,\n        \"count\": 1024\n      }\n    ],\n    \"spring\":
{\n      \"sids\": {\n        \"size\": 1000\n      }\n    },\n
\"route_reflector\": {\n      \"clusters\": [\n        {\n
\"cluster\": \"10.1.1.1\"\n      },\n      {\n
\": \"10.2.2.2\"\n    },\n      {\n
\"10.3.3.3\"\n    }\n    ]\n      }\n    }\n  }",
        "options": {
            "raw": {
                "language": "json"
            }
        }
    }
},
"url": {
    "raw": "https://{{server}}:{{port}}/service-orchestration/api/v1/orgs/
{{ORG}}/order",

```

```

        "protocol": "https",
        "host": [
            "{{server}}"
        ],
        "port": "{{port}}",
        "path": [
            "service-orchestration",
            "api",
            "v1",
            "orgs",
            "{{ORG}}",
            "order"
        ]
    },
    "response": []
},
{
    "name": "06-Exec Routing Resources",
    "event": [
        {
            "listen": "prerequisite",
            "script": {
                "exec": [
                    "var x=CryptoJS.enc.Utf8.parse(postman.getEnvironmentVariable(\"User
                    \")+\":\")+postman.getEnvironmentVariable(\"Password\"));";",
                    "var authHeader=CryptoJS.enc.Base64.stringify(x);",
                    "pm.request.headers.add({key: \"Authorization\", value: \"Basic
                    \"+authHeader});";",
                    ""
                ],
                "type": "text/javascript"
            }
        },
        {
            "listen": "test",
            "script": {
                "exec": [
                    ""
                ],
                "type": "text/javascript"
            }
        }
    ]
}

```

```

    ],
    "request": {
      "method": "POST",
      "header": [],
      "url": {
        "raw": "https://{{server}}:{{port}}/service-orchestration/api/v1/orgs/
{{ORG}}/order/customers/network-operator/instances/routing-stuff/exec",
        "protocol": "https",
        "host": [
          "{{server}}"
        ],
        "port": "{{port}}",
        "path": [
          "service-orchestration",
          "api",
          "v1",
          "orgs",
          "{{ORG}}",
          "order",
          "customers",
          "network-operator",
          "instances",
          "routing-stuff",
          "exec"
        ]
      }
    },
    "response": []
  },
  {
    "name": "07-Verify Resources",
    "event": [
      {
        "listen": "prerequest",
        "script": {
          "exec": [
            "var x=CryptoJS.enc.Utf8.parse(postman.getEnvironmentVariable(\"User
\")+\":\":"+postman.getEnvironmentVariable(\"Password\"));",
            "var authHeader=CryptoJS.enc.Base64.stringify(x);",
            "pm.request.headers.add({key: \"Authorization\", value: \"Basic
\"+authHeader});",
            ""
          ],
        }
      }
    ]
  }
],

```

```

        "type": "text/javascript"
    }
},
{
    "listen": "test",
    "script": {
        "exec": [
            ""
        ],
        "type": "text/javascript"
    }
}
],
"request": {
    "method": "GET",
    "header": [],
    "url": {
        "raw": "https://{{server}}:{{port}}/service-orchestration/api/v1/orgs/{{ORG}}/placement/network-elements",
        "protocol": "https",
        "host": [
            "{{server}}"
        ],
        "port": "{{port}}",
        "path": [
            "service-orchestration",
            "api",
            "v1",
            "orgs",
            "{{ORG}}",
            "placement",
            "network-elements"
        ]
    }
},
"response": []
}
]
}

```

Table 1 on page 692 lists the APIs in the sample Postman collection file for onboarding devices.

Table 160: REST APIs in the Sample Postman Collection File for Onboarding Devices.

REST API	Description	Reference in Collection File
Get Organization Details	Get credentials for accessing an organization and the organization details.	01-Who am I and get orgs
Get Site	Get the site where the device is to be installed and onboarded.	02-pick-site
Add L3 Address; see "Sample REST API to Create an IPv4 Address Pool" on page 692	Create layer 3 (L3) address groups.	03-Create L3 Addr
Post L3 Address	Save the L3 address groups in the database.	04-Exec L3 Addr
Add Routing Resources; see "Sample REST API to Create BGP Cluster ID Pool" on page 693	Create BGP cluster groups.	05-Create Routing Resources
Post Routing Resources	Save the BGP cluster groups in the database.	06-Exec Routing Resources
Get Resources	Get the L3 address groups and BGP clusters that were created for verification.	07-Verify Resources

Sample REST API to Create an IPv4 Address Pool



NOTE: The operation field in the JSON file can take up the following values:

- create—Creates new network resources if none exist. However, if resources already exist, new network resources specified in the JSON file are added to the existing ones.

- modify—Overrides the existing network resources with the values passed through the JSON file.
- delete—Removes the network resources specified in the JSON file.

The following is a sample of the REST API to create an IPv4 address resource pool:

```
https://{server}:{port}/service-orchestration/api/v1/orgs/{ORG}/order"
{
  "customer_id": "network-operator",
  "design_id": "l3-addr"
  "instance_id": "l3-addr",
  "operation": "create",
  "org_id": "<ORG>",
  "l3_addr": {
    "loopbacks": [{
      "name": "range-192",
      "prefix": "10.10.192.0/18"
    }]
    "ipv4_prefixes": [{
      "name": "pool-11",
      "prefix": "10.10.11.0/24"
    },
    {
      "name": "pool-12",
      "prefix": "10.10.12.0/24"
    }
  ]
}
}
```

Sample REST API to Create BGP Cluster ID Pool



NOTE: The operation field in the JSON file can take up the following values:

- create—Creates new network resources if none exist. However, if resources already exist, new network resources specified in the JSON file are added to the existing ones.

- **modify**—Overrides the existing network resources with the values passed through the JSON file.
- **delete**—Removes the network resources specified in the JSON file.

The following is a sample of the REST API to the create BGP cluster ID resource pool:

```
"https://{server}://{port}/service-orchestration/api/v1/orgs/{ORG}/order"
```

```
{
  "customer_id": "network-operator",
  "design_id": "routing",
  "instance_id": "routing-stuff",
  "operation": "create",
  "routing": {
    "autonomous_system": [
      {
        "name": 65200,
        "count": 1024
      }
    ],
    "spring": {
      "sids": {
        "size": 1000
      }
    },
    "route_reflector": {
      "clusters": [
        {
          "cluster": "192.168.1.1"
        },
        {
          "cluster": "192.168.2.2"
        },
        {
          "cluster": "192.168.3.3"
        }
      ]
    }
  }
}
```

Sample Postman Collection Files for Service Orchestration

The following is a sample of the Postman collection file to provision an L3VPN service in your network.

```
{
  "info": {
    "_postman_id": "1bdacde8-64c9-4aaa-ae31-325aef127e44",
    "name": "Service Orchestration - L3VPN",
    "schema": "https://schema.getpostman.com/json/collection/v2.1.0/collection.json",
    "_exporter_id": "829664"
  },
  "item": [
    {
      "name": "Create Topology - Step 1.1",
      "request": {
        "auth": {
          "type": "basic",
          "basic": [
            {
              "key": "password",
              "value": "{{Password}}",
              "type": "string"
            },
            {
              "key": "username",
              "value": "{{User}}",
              "type": "string"
            }
          ]
        },
        "method": "POST",
        "header": [],
        "body": {
          "mode": "raw",
          "raw": "{\n  \"customer_id\": \"l3vpn-topology-cid\", \n  \"design_id\":\n  \"topo\", \n  \"instance_id\": \"l3vpn-topology-iid\", \n  \"operation\": \"create\", \n\n  \"topo\": {\n    \"pop\": [\n      {\n        \"name\":\n        \"6745739c-50dc-40b6-8ba1-72683d199362\", \n        \"pe\": [\n          {\n            \"name\":\n            \"00000000-0000-0000-1000-8828fb0ef680\", \n            \"access\":\n            [\n              {\n                \"name\":\n                \"et-0/0/5\", \n                \"type\": \"ethernet\n
```



```

        "path": [
            "service-orchestration",
            "api",
            "v1",
            "orgs",
            "{{ORG}}",
            "order"
        ]
    },
    "response": []
},
{
    "name": "EXEC Topology Step 1.2",
    "request": {
        "auth": {
            "type": "basic",
            "basic": [
                {
                    "key": "password",
                    "value": "{{Password}}",
                    "type": "string"
                },
                {
                    "key": "username",
                    "value": "{{User}}",
                    "type": "string"
                }
            ]
        },
        "method": "POST",
        "header": [],
        "body": {
            "mode": "raw",
            "raw": ""
        },
        "url": {
            "raw": "https://{{server}}:{{port}}/service-orchestration/api/v1/orgs/{{ORG}}/order/customers/l3vpn-topology-cid/instances/l3vpn-topology-iid/exec",
            "protocol": "https",
            "host": [
                "{{server}}"
            ],

```

```

        "port": "{{port}}",
        "path": [
            "service-orchestration",
            "api",
            "v1",
            "orgs",
            "{{ORG}}",
            "order",
            "customers",
            "l3vpn-topology-cid",
            "instances",
            "l3vpn-topology-iid",
            "exec"
        ]
    },
    "response": []
},
{
    "name": "Create vpn resources Step 2.1",
    "request": {
        "auth": {
            "type": "basic",
            "basic": [
                {
                    "key": "password",
                    "value": "{{Password}}",
                    "type": "string"
                },
                {
                    "key": "username",
                    "value": "{{User}}",
                    "type": "string"
                }
            ]
        },
        "method": "POST",
        "header": [],
        "body": {
            "mode": "raw",
            "raw": "{\n  \"customer_id\": \"L3VPN-vpn\", \n  \"design_id\": \"vpn\n\", \n  \"instance_id\": \"vpn\", \n  \"operation\": \"create\", \n  \"vpn\": {\n    \"route_distinguisher\": [\n      {\n        \"count\": 1024, \n"

```

```

    {"name": 1234, "count": 1024, "route_target": [
      {"name": 1234, "count": 1024},
      {"name": 1235, "count": 1024}
    ]},
    {
      "options": {
        "raw": {
          "language": "json"
        }
      },
      "url": {
        "raw": "https://{{server}}:{{port}}/service-orchestration/api/v1/orgs/{{ORG}}/order",
        "protocol": "https",
        "host": [
          "{{server}}"
        ],
        "port": "{{port}}",
        "path": [
          "service-orchestration",
          "api",
          "v1",
          "orgs",
          "{{ORG}}",
          "order"
        ]
      },
      "response": []
    },
    {
      "name": "EXEC vpn resources Step 2.2",
      "request": {
        "auth": {
          "type": "basic",
          "basic": [
            {
              "key": "password",
              "value": "{{Password}}",
              "type": "string"
            }
          ]
        }
      }
    }
  ]
}

```

```

        "key": "username",
        "value": "{{User}}",
        "type": "string"
    }
]
},
"method": "POST",
"header": [],
"url": {
    "raw": "https://{{server}}:{{port}}/service-orchestration/api/v1/orgs/
{{ORG}}/order/customers/L3VPN-vpn/instances/vpn/exec",
    "protocol": "https",
    "host": [
        "{{server}}"
    ],
    "port": "{{port}}",
    "path": [
        "service-orchestration",
        "api",
        "v1",
        "orgs",
        "{{ORG}}",
        "order",
        "customers",
        "L3VPN-vpn",
        "instances",
        "vpn",
        "exec"
    ]
}
},
"response": []
},
{
    "name": "Check Placement Resources 3 Copy",
    "request": {
        "auth": {
            "type": "basic",
            "basic": [
                {
                    "key": "password",
                    "value": "{{Password}}",
                    "type": "string"
                }
            ]
        }
    }
}

```

```

        },
        {
            "key": "username",
            "value": "{{User}}",
            "type": "string"
        }
    ]
},
"method": "GET",
"header": [],
"url": {
    "raw": "https://{{server}}:{{port}}/service-orchestration/api/v1/orgs/
{{ORG}}/placement/network-elements",
    "protocol": "https",
    "host": [
        "{{server}}"
    ],
    "port": "{{port}}",
    "path": [
        "service-orchestration",
        "api",
        "v1",
        "orgs",
        "{{ORG}}",
        "placement",
        "network-elements"
    ]
}
},
"response": []
}
]
}

```

Table 161 on page 702 lists the APIs in the sample collection file.

Table 161: REST APIs in the Sample Postman Collection File for L3VPN Provisioning

REST API	Description	Reference in the Collection File
Upload topology resources service order	Create topology network resource pools for provisioning L3VPN service.	Create Topology - Step 1.1
Execute topology resources service order	Upload topology resource pools to the Routing Director database.	EXEC Topology Step 1.2
Upload VPN resources service order	Create VPN resource pools for provisioning L3VPN service.	Create vpn resources Step 2.1
Execute VPN resources service order	Upload VPN resource pools to the Routing Director database.	EXEC vpn resources Step 2.2
View network resources for placement	View available network resources to assign placement configurations for L3VPN service.	Check Placement Resources 3 Copy

The following is a sample of the Postman collection file to provision an L2VPN service in your network.

```
{
  "info": {
    "_postman_id": "dcbeaeaa-1cac-4edf-a601-a1836fb8b1af",
    "name": "L2VPN Documentation",
    "schema": "https://schema.getpostman.com/json/collection/v2.1.0/collection.json",
    "_exporter_id": "30272935"
  },
  "item": [
    {
      "name": "L2VPN S0 Upload",
      "request": {
        "auth": {
          "type": "basic",
          "basic": [
            {
              "key": "password",
              "value": "{{Password}}"
            }
          ]
        }
      }
    }
  ]
}
```



```

\", \n                                \"tagged_interface\":
{ \n                                \"type\": \"dot1q
\", \n                                \"dot1q_vlan_tagged\":
{ \n                                \"tg_type\": \"c-vlan
\", \n                                \"cvlan_id\":
200 \n                                }
\n                                }, \n
\"untagged_interface\": { \n                                \"speed\":
10000, \n                                \"lldp\":
true, \n                                \"oam_802.3ah_link\":
{ \n                                \"enabled\": true
\n                                } \n                                }
\n                                }, \n                                \"vpn_attachment\":
{ \n                                \"vpn_id\":
\n12vpn-1 \n                                }, \n                                \"service\":
{ \n                                \"svc_bandwidth\":
{ \n                                \"bandwidth\":
[ \n                                \"direction\": \"input-bw
\", \n                                \"cir\":
100000, \n                                \"cbs\":
100000 \n                                }, \n                                \"direction\": \"output-bw
\", \n                                \"cir\":
100000, \n                                \"cbs\":
100000 \n                                } \n                                ]
\n                                }
\n                                }, \n                                \"mac_addr_limit\":
{ \n                                \"limit_number\":
500, \n                                \"action\": \"drop\" \n                                }
\n                                } \n                                ] \n                                }
\n                                }, \n                                { \n                                \"site_id\":
\"PE2\", \n                                \"site_vpn_flavor\": \"site-vpn-flavor-single
\", \n                                \"devices\": { \n                                \"device\":
[ \n                                { \n                                \"device_id\":
\"ce2\", \n                                \"location\": \"BNG\" \n                                }
\n                                ] \n                                }, \n                                \"locations\":
{ \n                                \"location\": [ \n                                { \n                                \"location_id\": \"BNG\", \n                                \"address\": \"123 Bengaluru\", \n                                \"postal_code\":
\"20321\", \n                                \"state\": \"Karnataka
\", \n                                \"city\": \"Bengaluru\", \n                                \"country_code\": \"IN\" \n                                } \n                                ]

```



```

\n          }\n          ]\n          }\n          }\n\n          "options": {\n            "raw": {\n              "language": "json"\n            }\n          }\n        },\n        "url": {\n          "raw": "https://{{server}}:{{port}}/service-orchestration/api/v1/orgs/{{ORG}}/order",\n          "protocol": "https",\n          "host": [\n            "{{server}}"\n          ],\n          "port": "{{port}}",\n          "path": [\n            "service-orchestration",\n            "api",\n            "v1",\n            "orgs",\n            "{{ORG}}",\n            "order"\n          ]\n        }\n      },\n      "response": []\n    },\n    {\n      "name": "L2VPN S0 Exec",\n      "request": {\n        "auth": {\n          "type": "basic",\n          "basic": [\n            {\n              "key": "password",\n              "value": "{{Password}}",\n              "type": "string"\n            },\n            {\n              "key": "username",\n              "value": "{{User}}",\n              "type": "string"\n            }\n          ]\n        }\n      }\n    }\n  ]\n}

```

```

        ]
    },
    "method": "POST",
    "header": [],
    "body": {
        "mode": "raw",
        "raw": "",
        "options": {
            "raw": {
                "language": "json"
            }
        }
    },
    "url": {
        "raw": "https://{{server}}:{{port}}/service-orchestration/api/v1/orgs/
{{ORG}}/order/customers/l2vpn-so-cid/instances/l2vpn-so-iid/exec",
        "protocol": "https",
        "host": [
            "{{server}}"
        ],
        "port": "{{port}}",
        "path": [
            "service-orchestration",
            "api",
            "v1",
            "orgs",
            "{{ORG}}",
            "order",
            "customers",
            "l2vpn-so-cid",
            "instances",
            "l2vpn-so-iid",
            "exec"
        ]
    },
    "response": []
},
{
    "name": "VPN SO Upload",
    "request": {
        "auth": {
            "type": "basic",

```

```

        "basic": [
            {
                "key": "password",
                "value": "{{Password}}",
                "type": "string"
            },
            {
                "key": "username",
                "value": "{{User}}",
                "type": "string"
            }
        ],
        "method": "POST",
        "header": [],
        "body": {
            "mode": "raw",
            "raw": "{\n  \"customer_id\": \"l2vpn-so-cid\", \n  \"design_id\": \"vpn\n\", \n  \"instance_id\": \"l2vpn-vpn-iid\", \n  \"operation\": \"create\", \n  \"vpn\":\n{\n  \"route_distinguisher\": [\n    {\n      \"name\":\n1000, \n      \"count\": 1024\n    }, \n    {\n      \"name\":\n2000, \n      \"count\": 1024\n    }\n  ], \n  \"route_target\":\n[\n    {\n      \"name\": 500, \n      \"count\":\n1024\n    }, \n    {\n      \"name\": 600, \n      \"count\":\n1024\n    }\n  ]\n}\n}",
            "options": {
                "raw": {
                    "language": "json"
                }
            }
        },
        "url": {
            "raw": "https://{{server}}:{{port}}/service-orchestration/api/v1/orgs/\n{{ORG}}/order",
            "protocol": "https",
            "host": [
                "{{server}}"
            ],
            "port": "{{port}}",
            "path": [
                "service-orchestration",
                "api",
                "v1",

```

```

        "orgs",
        "{{ORG}}",
        "order"
    ]
}
},
"response": []
},
{
    "name": "VPN SO Exec",
    "request": {
        "auth": {
            "type": "basic",
            "basic": [
                {
                    "key": "password",
                    "value": "{{Password}}",
                    "type": "string"
                },
                {
                    "key": "username",
                    "value": "{{User}}",
                    "type": "string"
                }
            ]
        },
        "method": "POST",
        "header": [],
        "body": {
            "mode": "raw",
            "raw": "",
            "options": {
                "raw": {
                    "language": "json"
                }
            }
        },
        "url": {
            "raw": "https://{{server}}:{{port}}/service-orchestration/api/v1/orgs/
{{ORG}}/order/customers/l2vpn-so-cid/instances/l2vpn-vpn-iid/exec",
            "protocol": "https",
            "host": [
                "{{server}}"
            ]
        }
    }
}

```

```

        ],
        "port": "{{port}}",
        "path": [
            "service-orchestration",
            "api",
            "v1",
            "orgs",
            "{{ORG}}",
            "order",
            "customers",
            "l2vpn-so-cid",
            "instances",
            "l2vpn-vpn-iid",
            "exec"
        ]
    },
    "response": []
},
{
    "name": "L2 Resource Upload",
    "request": {
        "auth": {
            "type": "basic",
            "basic": [
                {
                    "key": "password",
                    "value": "{{Password}}",
                    "type": "string"
                },
                {
                    "key": "username",
                    "value": "{{User}}",
                    "type": "string"
                }
            ]
        },
        "method": "POST",
        "header": [],
        "body": {
            "mode": "raw",
            "raw": "{\n  \"customer_id\": \"l2vpn-so-cid\", \n  \"design_id\": \"l2-addr\", \n  \"instance_id\": \"l2-stuff\", \n  \"operation\": \"create\", \n  \"l2_addr\":

```

```

{\n      \lacp\": {\n          \system_id\": [\n          {\n
\name\": \"00:00:00:00:00\", \n          \count\":
254\n      }, \n          {\n          \name\":
\"00:00:00:00:01\", \n          \count\": 254\n      }
\n      ], \n          \admin_key\": [\n          {\n          \start
\": 200, \n          \size\": 10\n          }\n      ]
\n      }, \n          \evpn_esl\": [\n          {\n          \name\":
\"00:00:00:00:00:00:00:00\", \n          \count\": 254\n      }, \n
{\n          \name\": \"00:00:00:00:00:00:00:01\", \n          \count\":
254\n      }\n      ]\n      }\n      },
      \"options\": {
        \"raw\": {
          \"language\": \"json\"
        }
      },
      \"url\": {
        \"raw\": \"https://{{server}}:{{port}}/service-orchestration/api/v1/orgs/
{{ORG}}/order\",
        \"protocol\": \"https\",
        \"host\": [
          \"{{server}}\"
        ],
        \"port\": \"{{port}}\",
        \"path\": [
          \"service-orchestration\",
          \"api\",
          \"v1\",
          \"orgs\",
          \"{{ORG}}\",
          \"order\"
        ]
      }
    },
    \"response\": []
  },
  {
    \"name\": \"L2 Resource Exec\",
    \"request\": {
      \"auth\": {
        \"type\": \"basic\",
        \"basic\": [
          {

```



```

        "key": "password",
        "value": "{{Password}}",
        "type": "string"
    },
    {
        "key": "username",
        "value": "{{User}}",
        "type": "string"
    }
]
},
"method": "POST",
"header": [],
"body": {
    "mode": "raw",
    "raw": "",
    "options": {
        "raw": {
            "language": "json"
        }
    }
},
"url": {
    "raw": "https://{{server}}:{{port}}/service-orchestration/api/v1/orgs/
{{ORG}}/order/customers/l2vpn-so-cid/instances/l2-stuff/exec",
    "protocol": "https",
    "host": [
        "{{server}}"
    ],
    "port": "{{port}}",
    "path": [
        "service-orchestration",
        "api",
        "v1",
        "orgs",
        "{{ORG}}",
        "order",
        "customers",
        "l2vpn-so-cid",
        "instances",
        "l2-stuff",
        "exec"
    ]
}

```

```

    }
  },
  "response": []
},
{
  "name": "Topology SO Upload",
  "request": {
    "auth": {
      "type": "basic",
      "basic": [
        {
          "key": "password",
          "value": "{{Password}}",
          "type": "string"
        },
        {
          "key": "username",
          "value": "{{User}}",
          "type": "string"
        }
      ]
    },
    "method": "POST",
    "header": [],
    "body": {
      "mode": "raw",
      "raw": "{\n  \"customer_id\": \"l2vpn-so-cid\",\n  \"design_id\": \"topo\n\n  \"instance_id\": \"l2vpn-topo-iid\",\n  \"operation\": \"create\",\n  \"topo\":\n{\n  \"pop\": [\n    {\n      \"name\": \"76025add-2f0a-4959-857f-a323c9a4e2d2\",\n      \"pe\": [\n        {\n          \"name\n\n      \": \"00000000-0000-0000-1000-74e798609000\",\n          \"access\":\n[\n        {\n          \"name\":\n\n      \"et-0/0/14\",\n          \"type\": \"ethernet\n\n      \",\n          \"speed\": 10000,\n          \"ce\":\n\n      \"ce1\",\n          \"vlans\": {\n\n      \"start\": 200,\n          \"size\":\n\n      10\n        }\n      }\n    ],\n    \"bandwidth\":\n\n      \"2000000\",\n    \"routes\": \"1000\",\n    \"mac_addrs\n\n      \": \"1000\"\n    },\n    \"postal_code_matches\n\n      \": [\n    {\n      \"name\": \"SVL\n\n      \",\n      \"regex\": \"10...\"\n    }\n  ],\n  }\n},\n  {\n    \"name\":

```

```

\ "9070a569-3650-4b50-9451-5acdbd8ccf41", \n                \ "pe": [ \n
{ \n                \ "name": \ "00000000-0000-0000-1000-
d4996c5731e8", \n                \ "access": [ \n
{ \n                \ "name": \ "xe-1/0/3", \n
\ "type": \ "ethernet", \n                \ "speed":
10000, \n                \ "ce": \ "ce2", \n                \ "vlans
": { \n                \ "start": 200, \n
\ "size": 10 \n                } \n                }
\ n                ], \n                \ "bandwidth":
\ "100000", \n                \ "routes": \ "1000", \n                \ "mac_addrs
": \ "1000" \n                } \n                \ n                ], \n
\ "postal_code_matches": [ \n                { \n                \ "name": \ "BNG
", \n                \ "regex": \ "20..." \n                } \n                ]
\ n                } \n                ] \n                } \n",
                "options": {
                    "raw": {
                        "language": "json"
                    }
                },
                "url": {
                    "raw": "https://{{server}}:{{port}}/service-orchestration/api/v1/orgs/
{{ORG}}/order",
                    "protocol": "https",
                    "host": [
                        "{{server}}"
                    ],
                    "port": "{{port}}",
                    "path": [
                        "service-orchestration",
                        "api",
                        "v1",
                        "orgs",
                        "{{ORG}}",
                        "order"
                    ]
                }
            },
            "response": []
        },
        {
            "name": "Topology SO Exec",
            "request": {

```

```

    "auth": {
      "type": "basic",
      "basic": [
        {
          "key": "password",
          "value": "{{Password}}",
          "type": "string"
        },
        {
          "key": "username",
          "value": "{{User}}",
          "type": "string"
        }
      ]
    },
    "method": "POST",
    "header": [],
    "body": {
      "mode": "raw",
      "raw": "",
      "options": {
        "raw": {
          "language": "json"
        }
      }
    },
    "url": {
      "raw": "https://{{server}}:{{port}}/service-orchestration/api/v1/orgs/
{{ORG}}/order/customers/l2vpn-so-cid/instances/l2vpn-topo-iid/exec",
      "protocol": "https",
      "host": [
        "{{server}}"
      ],
      "port": "{{port}}",
      "path": [
        "service-orchestration",
        "api",
        "v1",
        "orgs",
        "{{ORG}}",
        "order",
        "customers",
        "l2vpn-so-cid",

```

```

        "instances",
        "l2vpn-topo-iid",
        "exec"
    ]
  },
},
"response": []
}
]
}

```

Table 162 on page 716 lists the APIs in the sample collection file.

Table 162: REST APIs in the Sample Postman Collection File for L2VPN Provisioning

REST API	Description	Reference in the Collection File
Upload L2VPN service order	Upload or save the L2VPN service order to the Routing Director database.	L2VPN SO Upload
Execute the L2VPN service order	Activate the provisioning workflow for the L2VPN service order.	L2VPN SO Exec
Upload VPN resources service order	Create VPN resource pools for provisioning an L2VPN service.	VPN SO Upload
Execute VPN resources service order	Upload VPN resource pools to the Routing Director database. These resources are allocated when provisioning an L2VPN service.	VPN SO Exec
Upload L2 address resources service order	Create L2 address resource pools for provisioning an L2VPN service.	L2 Resource Upload
Execute L2 address resources service order	Upload L2 address resource pools to the Routing Director database. These resources are allocated when provisioning an L2VPN service.	L2 Resource Exec

Table 162: REST APIs in the Sample Postman Collection File for L2VPN Provisioning (Continued)

REST API	Description	Reference in the Collection File
Upload topology resources service order	Create topology resource pools for provisioning an L2VPN service.	Topology SO Upload
Execute topology resources service order	Upload topology resource pools to the Routing Director database. These resources are allocated when provisioning an L2VPN service.	Topology SO Exec

The following is a sample of the Postman collection file to provision an L2 circuit service in your network.

```
{
  "info": {
    "_postman_id": "f5eabd22-790c-4cfc-9f28-6f7a35644b06",
    "name": "L2VPN NSM",
    "schema": "https://schema.getpostman.com/json/collection/v2.1.0/collection.json",
    "_exporter_id": "16733671"
  },
  "item": [
    {
      "name": "Create L2VPN Instance",
      "request": {
        "auth": {
          "type": "basic",
          "basic": [
            {
              "key": "password",
              "value": "{{password}}",
              "type": "string"
            },
            {
              "key": "username",
              "value": "{{user}}",
              "type": "string"
            }
          ]
        }
      }
    }
  ]
},
```



```

{"service": {
  "svc_input_bandwidth": {
    "input_bandwidth": {
      "cbs": {
        "cir": {
          "type": "bw-per-port"
        }
      }
    }
  },
  "svc_output_bandwidth": {
    "output_bandwidth": {
      "cbs": {
        "cir": {
          "type": "bw-per-port"
        }
      }
    }
  }
},
"vpn_node_id": "yang-node-2",
"00000000-0000-0000-1000-407f5f32f140",
{
  "interface": {
    "et-0/0/22:0",
    "vlan": {
      "agent": {
        "interface": {
          "signaling_options": {
            "type": {
              "t_ldp_pwe_type": "vpws-type"
            },
            "ac_pw_list": {
              "peer_addr": {
                "vc_id": {
                  "id": {
                    "description": "CE2 description"
                  },
                  "port_id": {
                    "connection": {
                      "encapsulation":

```



```

        "api",
        "v1",
        "orgs",
        "{{ORG}}",
        "order"
    ]
}
},
"response": []
},
{
    "name": "Exec L2VPN Instance",
    "request": {
        "auth": {
            "type": "basic",
            "basic": [
                {
                    "key": "password",
                    "value": "{{password}}",
                    "type": "string"
                },
                {
                    "key": "username",
                    "value": "{{user}}",
                    "type": "string"
                }
            ]
        },
        "method": "POST",
        "header": [],
        "body": {
            "mode": "raw",
            "raw": "",
            "options": {
                "raw": {
                    "language": "json"
                }
            }
        },
        "url": {
            "raw": "https://{{server}}:{{port}}/service-orchestration/api/v1/orgs/
{{ORG}}/order/customers/Widget/instances/Widget/exec",
            "protocol": "https",

```

```

        "host": [
            "{{server}}"
        ],
        "port": "{{port}}",
        "path": [
            "service-orchestration",
            "api",
            "v1",
            "orgs",
            "{{ORG}}",
            "order",
            "customers",
            "Widget",
            "instances",
            "Widget",
            "exec"
        ]
    }
},
"response": []
},
{
    "name": "Get L2VPN Instance",
    "protocolProfileBehavior": {
        "disableBodyPruning": true
    },
    "request": {
        "auth": {
            "type": "basic",
            "basic": [
                {
                    "key": "password",
                    "value": "{{password}}",
                    "type": "string"
                },
                {
                    "key": "username",
                    "value": "{{user}}",
                    "type": "string"
                }
            ]
        },
        "method": "GET",

```

```

        "header": [],
        "body": {
            "mode": "raw",
            "raw": "",
            "options": {
                "raw": {
                    "language": "json"
                }
            }
        },
    },
    "url": {
        "raw": "https://{{server}}:{{port}}/service-orchestration/api/v1/orgs/{{ORG}}/order/customers/Widget/instances/Widget",
        "protocol": "https",
        "host": [
            "{{server}}"
        ],
        "port": "{{port}}",
        "path": [
            "service-orchestration",
            "api",
            "v1",
            "orgs",
            "{{ORG}}",
            "order",
            "customers",
            "Widget",
            "instances",
            "Widget"
        ]
    }
},
"response": []
},
{
    "name": "Delete L2VPN Instance",
    "request": {
        "auth": {
            "type": "basic",
            "basic": [
                {
                    "key": "password",
                    "value": "{{password}}",

```

```

        "type": "string"
    },
    {
        "key": "username",
        "value": "{{user}}",
        "type": "string"
    }
]
},
"method": "POST",
"header": [],
"body": {
    "mode": "raw",
    "raw": "{\n  \"customer_id\": \"Widget\", \n  \"design_id\": \"l2vpn-nsm\n\", \n  \"instance_id\": \"Widget\", \n  \"l2vpn_ntw\": {}, \n  \"fh_config\": {\n    \"type\":\n    \"service_order\"\n  }, \n  \"operation\": \"delete\"\n}",
    "options": {
        "raw": {
            "language": "json"
        }
    }
},
"url": {
    "raw": "https://{{server}}:{{port}}/service-orchestration/api/v1/orgs/{{ORG}}/order",
    "protocol": "https",
    "host": [
        "{{server}}"
    ],
    "port": "{{port}}",
    "path": [
        "service-orchestration",
        "api",
        "v1",
        "orgs",
        "{{ORG}}",
        "order"
    ]
}
},
"response": []
},
{

```

```
"name": "Create VPN Resources",
"event": [
  {
    "listen": "prerequisite",
    "script": {
      "exec": [
        "var x=CryptoJS.enc.Utf8.parse(postman.getEnvironmentVariable(\"User\\")+\"\\\\\":\"+postman.getEnvironmentVariable(\"Password\\\\\"));\",
        \"var authHeader=CryptoJS.enc.Base64.stringify(x);\",
        \"pm.request.headers.add({key: \\\"Authorization\\\", value: \\\"Basic\\\"+authHeader});\"
      ],
      "type": "text/javascript"
    }
  }
],
"request": {
  "auth": {
    "type": "noauth"
  },
  "method": "POST",
  "header": [],
  "body": {
    "mode": "raw",
    "raw": "{\n  \\"customer_id\\": \\"network-operator\\",\n  \\"design_id\\": \\"vpn\\",\n  \\"fh_config\\": {\n    \\"type\\": \\"service_order\\\"\n  },\n  \\"instance_id\\": \\"rds-and-rtss\\",\n  \\"operation\\": \\"create\\",\n  \\"vpn\\": {\n    \\"route_distinguisher\\": [\n      {\n        \\"count\\": 1024,\n        \\"name\\": 1234\n      },\n      {\n        \\"count\\": 1024,\n        \\"name\\": 1235\n      }\n    ],\n    \\"route_target\\": [\n      {\n        \\"count\\": 1024,\n        \\"name\\": 1234\n      },\n      {\n        \\"count\\": 1024,\n        \\"name\\": 1235\n      }\n    ],\n    \\"community\\": [\n      {\n        \\"name\\": \\"target:64500\\",\n        \\"count\\": 10000\n      }\n    ],\n    \\"virtual_circuit_id\\": [\n      {\n        \\"start\\": 1,\n        \\"count\\": 10000\n      }\n    ]\n  }\n}",
    "options": {
      "raw": {
        "language": "json"
      }
    }
  }
},
"url": {
```

```

        "raw": "https://{{server}}:{{port}}/service-orchestration/api/v1/orgs/
{{ORG}}/order",
        "protocol": "https",
        "host": [
            "{{server}}"
        ],
        "port": "{{port}}",
        "path": [
            "service-orchestration",
            "api",
            "v1",
            "orgs",
            "{{ORG}}",
            "order"
        ]
    },
    "response": []
},
{
    "name": "Exec VPN Resources",
    "request": {
        "auth": {
            "type": "basic",
            "basic": [
                {
                    "key": "password",
                    "value": "{{password}}",
                    "type": "string"
                },
                {
                    "key": "username",
                    "value": "{{user}}",
                    "type": "string"
                }
            ]
        },
        "method": "POST",
        "header": [],
        "body": {
            "mode": "raw",
            "raw": "{\n  \"customer_id\": \"internal\", \n  \"design_id\": \"l2vpn\", \n
            \"instance_id\": \"l2vpn\", \n  \"l2vpn_svc\": {\n    \"sites\": {\n      \"site\": [\n

```

```

{
  \n
  \"devices\": {
    \n
    \"device\": [
      \n
      {
        \n
        \"device_id\": \"CE4\",
        \n
        \"location\": \"LAX\"
      }
    ]
  },
  \n
  \"locations\": {
    \n
    \"location\": [
      \n
      {
        \n
        \"address\": \"123 Rodeo Drive\",
        \n
        \"city\": \"Beverly Hills\",
        \n
        \"country_code\": \"US\",
        \n
        \"location_id\": \"LAX\",
        \n
        \"postal_code\": \"90210\",
        \n
        \"state\": \"CA\"
      }
    ]
  },
  \n
  \"mac_loop_prevention\": {
    \n
    \"frequency\": 6,
    \n
    \"retry_timer\": 10
  },
  \n
  \"site_id\": \"Head Office\",
  \n
  \"site_network_accesses\": {
    \n
    \"site_network_access\": [
      \n
      {
        \n
        \"access_diversity\": {
          \n
          \"constraints\": {
            \n
            \"constraint\": {
              \n
              \"constraint_type\": \"pe-diverse\",
              \n
              \"target\": {
                \n
                \"group\": {
                  \n
                  \"group_id\": \"CE4-DG\"
                }
              }
            }
          }
        }
      }
    ]
  },
  \n
  \"groups\": {
    \n
    \"group\": [
      \n
      {
        \n
        \"group_id\": \"CE4-DG\"
      }
    ]
  },
  \n
  \"availability\": {
    \n
    \"all_active\": [
      \n
      null
    ]
  },
  \n
  \"bearer\": {
    \n
    \"requested_type\": {
      \n
      \"strict\": true,
      \n
      \"type\": \"ethernet\"
    }
  },
  \n
  \"connection\": {
    \n
    \"eth_inf_type\": \"tagged\",
    \n
    \"lag_interfaces\": {
      \n
      \"lag_interface\": [
        \n
        {
          \n
          \"index\": \"CE4-LAG1\",
          \n
          \"lcp\": {
            \n
            \"enabled\": true,
            \n
            \"member_links\": {
              \n
              \"member_link\": [
                \n
                {
                  \n
                  \"name\": \"CE4-LAG1-link1\",
                  \n
                  \"oam_802.3ah_link\": {
                    \n
                    \"enabled\": true
                  }
                }
              ]
            }
          }
        }
      ]
    }
  },
  \n
  \"tagged_interface\": {
    \n
    \"dot1q_vlan_tagged\": {
      \n
      \"cvlan_id\": 200,
      \n
      \"tg_type\": \"c-vlan\"
    },
    \n
    \"type\": \"dot1q\"
  },
  \n
  \"device_reference\": {
    \n
    \"mac_addr_limit\": {
      \n
      \"action\": \"drop\",
      \n
      \"limit_number\": 500
    }
  },
  \n
  \"network_access_id\": \"CE4-1\",
  \n
  \"service\": {
    \n
    \"svc_bandwidth\": {
      \n
      \"bandwidth\": [
        \n
        {
          \n
          \"cbs\": \"10000\",
          \n
          \"cir\": \"700000\",
          \n
          \"direction\": \"input-bw\"
        }
      ]
    }
  }
}

```



```

\ " \n
\ "10000\ " , \n
\ "output-bw\ " \n
\ n
\ n
\ "Sales-EVPN\ " \n
\ "access_diversity\ ": { \n
\ ": [ \n
\ " , \n
[ \n
\ " \n
\ n
\ n
\ "groups\ ": { \n
{ \n
\ n
\ n
\ ": { \n
\ n
\ n
{ \n
\ " \n
\ " \n
\ " \n
{ \n
\ "index\ ": \ "CE4-LAG1\ " , \n
\ "enabled\ ": true , \n
\ "member_link\ ": [ \n
\ "CE4-LAG1-link2\ " , \n
{ \n
\ n
\ n
40000 \n
\ n
\ n
\ n
{ \n
200 , \n
\ "type\ ": \ "dot1q\ " \n
\ ": \ "CE4\ " , \n
\ " , \n
\ "network_access_id\ ": \ "CE4-2\ " , \n
\ "svc_bandwidth\ ": { \n
{ \n
\ "70000\ " , \n
\ " \n
\ "10000\ " , \n
\ "output-bw\ " \n
\ n
\ "Sales-EVPN\ " \n

```

```

\"site_vpn_flavor\": \"site-vpn-flavor-single\"\\n      },\\n      {\\n          \"devices\":
{\\n          \"device\": [\\n          {\\n          \"device_id\":
\\\"CE6\\\",\\n          \"location\": \"LAX\"\\n          }\\n          ]
\\n      },\\n      \"locations\": {\\n          \"location\": [\\n
{\\n          \"address\": \"123 Rodeo Drive\",\\n          \"city\": \"Beverly Hills
\\\",\\n          \"country_code\": \"US\",\\n          \"location_id\": \"LAX
\\\",\\n          \"postal_code\": \"90210\",\\n          \"state\": \"CA
\\\"\\n          }\\n          ]\\n          },\\n          \"mac_loop_prevention\":
{\\n          \"frequency\": 6,\\n          \"retry_timer\": 10\\n          },\\n
\\\"site_id\": \"LAX depot\",\\n          \"site_network_accesses\": {\\n
\\\"site_network_access\": [\\n          {\\n          \"bearer\": {\\n
\\\"requested_type\": {\\n          \"strict\": true,\\n          \"type\":
\\\"ethernet\"\\n          }\\n          },\\n          \"connection\":
{\\n          \"eth_inf_type\": \"tagged\",\\n          \"tagged_interface\":
{\\n          \"dot1q_vlan_tagged\": {\\n          \"cvlan_id\":
102,\\n          \"tg_type\": \"c-vlan\"\\n          },\\n
\\\"type\": \"dot1q\"\\n          },\\n          \"untagged_interface\":
{\\n          \"lldp\": true,\\n          \"oam_802.3ah_link\":
{\\n          \"enabled\": true\\n          },\\n          \"speed
\\\": 40000\\n          }\\n          },\\n          \"device_reference\":
\\\"CE6\\\",\\n          \"mac_addr_limit\": {\\n          \"action\": \"drop
\\\",\\n          \"limit_number\": 500\\n          },\\n
\\\"network_access_id\": \"CE6-1\",\\n          \"service\": {\\n
\\\"svc_bandwidth\": {\\n          \"bandwidth\": [\\n
{\\n          \"cbs\": \"10000\",\\n          \"cir\":
\\\"70000\",\\n          \"direction\": \"input-bw
\\\"\\n          },\\n          {\\n          \"cbs\":
\\\"10000\",\\n          \"cir\": \"50000\",\\n          \"direction\":
\\\"output-bw\"\\n          }\\n          ]\\n          }
\\n          },\\n          \"vpn_attachment\": {\\n          \"vpn_id\":
\\\"Sales-EVPN\"\\n          }\\n          }\\n          ]\\n          },\\n
\\\"site_vpn_flavor\": \"site-vpn-flavor-single\"\\n      },\\n      {\\n          \"devices\":
{\\n          \"device\": [\\n          {\\n          \"device_id\":
\\\"CE2\\\",\\n          \"location\": \"NYC1\"\\n          }\\n          ]
\\n      },\\n      \"locations\": {\\n          \"location\": [\\n
{\\n          \"address\": \"2922 Broadway Ave\",\\n          \"city\": \"New York City
\\\",\\n          \"country_code\": \"US\",\\n          \"location_id\":
\\\"NYC1\",\\n          \"postal_code\": \"10001\",\\n          \"state\": \"NY
\\\"\\n          }\\n          ]\\n          },\\n          \"mac_loop_prevention\":
{\\n          \"frequency\": 6,\\n          \"retry_timer\": 10\\n          },\\n
\\\"site_id\": \"Manufacturing\",\\n          \"site_network_accesses\": {\\n
\\\"site_network_access\": [\\n          {\\n          \"access_diversity\":
{\\n          \"constraints\": {\\n          \"constraint\":

```



```

        }
    }
},
"url": {
    "raw": "https://{{server}}:{{port}}/service-orchestration/api/v1/orgs/
{{ORG}}/order/customers/network-operator/instances/rds-and-rtts/exec",
    "protocol": "https",
    "host": [
        "{{server}}"
    ],
    "port": "{{port}}",
    "path": [
        "service-orchestration",
        "api",
        "v1",
        "orgs",
        "{{ORG}}",
        "order",
        "customers",
        "network-operator",
        "instances",
        "rds-and-rtts",
        "exec"
    ]
},
"response": []
},
{
    "name": "Get VPN Resources",
    "protocolProfileBehavior": {
        "disableBodyPruning": true
    },
    "request": {
        "auth": {
            "type": "basic",
            "basic": [
                {
                    "key": "password",
                    "value": "{{password}}",
                    "type": "string"
                }
            ],
            {

```

```

        "key": "username",
        "value": "{{user}}",
        "type": "string"
    }
]
},
"method": "GET",
"header": [],
"body": {
    "mode": "raw",
    "raw": "{\n  \"customer_id\": \"internal\", \n  \"design_id\": \"l2vpn\", \n
  \"instance_id\": \"l2vpn\", \n  \"l2vpn_svc\": {\n    \"sites\": {\n      \"site\": [\n
{ \n        \"devices\": {\n          \"device\": [\n            {\n
  \"device_id\": \"CE4\", \n          \"location\": \"LAX\" \n        } \n      ]
  \n    }, \n    \"locations\": {\n      \"location\": [\n
{ \n        \"address\": \"123 Rodeo Drive\", \n        \"city\": \"Beverly Hills
\", \n        \"country_code\": \"US\", \n        \"location_id\": \"LAX
\", \n        \"postal_code\": \"90210\", \n        \"state\": \"CA
\" \n      } \n    ] \n    }, \n    \"mac_loop_prevention\":
{ \n      \"frequency\": 6, \n      \"retry_timer\": 10 \n    }, \n
  \"site_id\": \"Head Office\", \n    \"site_network_accesses\": {\n
  \"site_network_access\": [\n      {\n        \"access_diversity\":
{ \n          \"constraints\": {\n            \"constraint\":
[ \n              {\n                \"constraint_type\": \"pe-diverse
\", \n              {\n                \"target\": {\n                  \"group\":
[ \n                    {\n                      \"group_id\": \"CE4-DG
\" \n                  } \n                } \n              ] \n            }
  \n          } \n        } \n      ], \n    }, \n
  \"groups\": {\n      \"group\": [\n
{ \n        \"group_id\": \"CE4-DG\" \n      }
  \n    ] \n    }, \n    \"availability
\": {\n      \"all_active\": [\n        null \n      ]
  \n    }, \n    \"bearer\": {\n      \"requested_type\":
{ \n        \"strict\": true, \n        \"type\": \"ethernet
\" \n      } \n    }, \n    \"connection\":
{ \n      \"eth_inf_type\": \"tagged\", \n      \"lag_interfaces\":
{ \n        \"lag_interface\": [\n          {\n
  \"index\": \"CE4-LAG1\", \n          \"lacp\": {\n
  \"enabled\": true, \n          \"member_links\": {\n
  \"member_link\": [\n            {\n              \"name\":
  \"CE4-LAG1-link1\", \n            \"oam_802.3ah_link\":
{ \n              \"enabled\": true
  \n            }, \n          \"speed\":

```

[illegible]

```

\"type\": \"dot1q\"\\n
\\\": \"CE4\",\\n
\\\",\\n
\\\"network_access_id\\\": \"CE4-2\",\\n
\\\"svc_bandwidth\\\": {\\n
{\\n
\\\"cbs\\\": \"10000\",\\n
\\\"700000\",\\n
\\\"\\n
\\\",\\n
\\\"10000\",\\n
\\\"output-bw\"\\n
\\n
\\\",\\n
\\\"Sales-EVPN\"\\n
\\\"site_vpn_flavor\\\": \"site-vpn-flavor-single\"\\n
{\\n
\\\"device\\\": [\\n
\\\"CE6\",\\n
\\n
\\\",\\n
{\\n
\\\"address\\\": \"123 Rodeo Drive\",\\n
\\\",\\n
\\\"country_code\\\": \"US\",\\n
\\\",\\n
\\\"postal_code\\\": \"90210\",\\n
\\\"\\n
\\\",\\n
{\\n
\\\"frequency\\\": 6,\\n
\\\"site_id\\\": \"LAX depot\",\\n
\\\"site_network_access\\\": [\\n
\\\"requested_type\\\": {\\n
\\\"ethernet\"\\n
{\\n
\\\"eth_inf_type\\\": \"tagged\",\\n
{\\n
\\\"dot1q_vlan_tagged\\\": {\\n
102,\\n
\\\"type\\\": \"dot1q\"\\n
{\\n
\\\"lldp\\\": true,\\n
{\\n
\\\"enabled\\\": true\\n
\\\": 40000\\n
\\\"CE6\",\\n
\\\",\\n
\\\"network_access_id\\\": \"CE6-1\",\\n
\\\"svc_bandwidth\\\": {\\n
{\\n
\\\"cbs\\\": \"10000\",\\n
\\\"700000\",\\n
\\\"\\n
\\\",\\n
\\\"10000\",\\n
\\\"output-bw\"\\n
\\n
\\\",\\n
\\\"Sales-EVPN\"\\n

```



```

\"site_vpn_flavor\": \"site-vpn-flavor-single\"\\n      },\\n      {\\n          \"devices\":
{\\n          \"device\": [\\n          {\\n          \"device_id\":
\\\"CE2\\\",\\n          \"location\": \"NYC1\"\\n          }\\n          ]
\\n      },\\n      \"locations\": {\\n          \"location\": [\\n
{\\n          \"address\": \"2922 Broadway Ave\",\\n          \"city\": \"New York City
\\\",\\n          \"country_code\": \"US\",\\n          \"location_id\":
\\\"NYC1\\\",\\n          \"postal_code\": \"10001\",\\n          \"state\": \"NY
\\\"\\n          }\\n          ]\\n          },\\n          \"mac_loop_prevention\":
{\\n          \"frequency\": 6,\\n          \"retry_timer\": 10\\n          },\\n
\\\"site_id\": \"Manufacturing\",\\n          \"site_network_accesses\": {\\n
\\\"site_network_access\": [\\n          {\\n          \"access_diversity\":
{\\n          \"constraints\": {\\n          \"constraint\":
[\\n          {\\n          \"constraint_type\": \"pe-diverse
\\\",\\n          \"target\": {\\n          \"group\":
[\\n          {\\n          \"group_id\": \"CE2-DG
\\\"\\n          }\\n          ]\\n          },\\n          }
\\n          }\\n          ]\\n          },\\n
\\\"groups\": {\\n          \"group\": [\\n
{\\n          \"group_id\": \"CE2-DG\"\\n          }
\\n          ]\\n          }\\n          },\\n          \"availability
\\\": {\\n          \"access_priority\": 1000,\\n          \"single_active\":
[\\n          null\\n          ]\\n          },\\n          \"bearer
\\\": {\\n          \"requested_type\": {\\n          \"strict\":
true,\\n          \"type\": \"ethernet\"\\n          }
\\n          },\\n          \"connection\": {\\n          \"eth_inf_type\":
\\\"tagged\\\",\\n          \"lag_interfaces\": {\\n          \"lag_interface\":
[\\n          {\\n          \"index\": \"CE2-
LAG1\",\\n          \"lACP\": {\\n          \"enabled\":
true,\\n          \"lldp\": true,\\n          \"member_links\":
{\\n          \"member_link\": [\\n
{\\n          \"name\": \"CE2-LAG1-
link1\",\\n          \"oam_802.3ah_link\":
{\\n          \"enabled\": true
\\n          },\\n          \"speed\":
40000\\n          }\\n          ]
\\n          }\\n          }\\n          }\\n          }
\\n          ],\\n          },\\n          \"tagged_interface\":
{\\n          \"dot1q_vlan_tagged\": {\\n          \"cvlan_id\":
100,\\n          \"tg_type\": \"c-vlan\"\\n          },\\n
\\\"type\": \"dot1q\"\\n          }\\n          },\\n          \"device_reference
\\\": \"CE2\",\\n          \"mac_addr_limit\": {\\n          \"action\": \"drop
\\\",\\n          \"limit_number\": 500\\n          },\\n
\\\"network_access_id\": \"CE2-1\",\\n          \"service\": {\\n

```



```

\"10000\", \n                                \"cir\": \"500000\", \n                                \"direction\":
\"output-bw\" \n                                } \n                                ] \n                                }
\n                                }, \n                                \"vpn_attachment\": { \n                                \"vpn_id\":
\"Sales-EVPN\" \n                                } \n                                ] \n                                }, \n
\"site_vpn_flavor\": \"site-vpn-flavor-single\" \n                                } \n                                ] \n                                }, \n                                \"vpn_services
\": { \n                                \"vpn_service\": [ \n                                { \n                                \"customer_name\": \"Widget\", \n
\"svc_topo\": \"any-to-any\", \n                                \"vpn_id\": \"Sales-EVPN\", \n                                \"vpn_svc_type
\": \"pbb-evpn\" \n                                } \n                                ] \n                                } \n                                }, \n                                \"fh_config\": { \n                                \"type\":
\"service_order\" \n                                }, \n                                \"operation\": \"delete\" \n                                },
                                \"options\": {
                                    \"raw\": {
                                        \"language\": \"json\"
                                    }
                                }
                            },
                            \"url\": {
                                \"raw\": \"https://{{server}}:{{port}}/service-orchestration/api/v1/orgs/
{{ORG}}/order/customers/Widget/instances/Widget\",
                                \"protocol\": \"https\",
                                \"host\": [
                                    \"{{server}}\"
                                ],
                                \"port\": \"{{port}}\",
                                \"path\": [
                                    \"service-orchestration\",
                                    \"api\",
                                    \"v1\",
                                    \"orgs\",
                                    \"{{ORG}}\",
                                    \"order\",
                                    \"customers\",
                                    \"Widget\",
                                    \"instances\",
                                    \"Widget\"
                                ]
                            }
                        },
                        \"response\": []
                    }
                ]
            }

```

Table 163 on page 739 lists the APIs in the sample collection file.

Table 163: REST APIs in the Sample Postman Collection File for L2 Circuit Provisioning

REST API	Description	Reference in Collection File
Create L2 circuit service instance	Create an L2 circuit service instance.	Create L2VPN Instance
Execute L2 circuit service instance	Activate the provisioning workflow for the L2 circuit service instance.	Exec L2VPN Instance
View L2 circuit provisioning status	View the status of provisioning of the L2 circuit service instance.	Get L2VPN Instance
Delete L2 circuit service instance	Activate the delete workflow for the L2 circuit service instance.	Delete L2VPN Instance
Create VPN resource pool	Create VPN resource pool for provisioning an L2 circuit service.	Create VPN Resources
Execute VPN resource instance	Upload VPN resource pool to the Routing Director database. These resources are allocated when provisioning an L2 circuit service.	Exec VPN Resources
Get VPN resources for L2 circuit service provisioning	View resources available in the VPN resource pool for provisioning an L2 circuit service.	Get VPN Resources

Sample REST API to Create Topology Resources for L3VPN and L2VPN Services

NOTE: The operation field in the JSON file can take up the following values:

- create—Creates new network resources if none exist. However, if resources already exist, new network resources specified in the JSON file are added to the existing ones.

- modify—Overrides the existing network resources with the values passed through the JSON file.
- delete—Removes the network resources specified in the JSON file.

The following is a sample of the REST API to create topology resource pool for L3VPN and L2VPN services:

```
https://{server}:{port}/service-orchestration/api/v1/orgs/{ORG}/order

# This file is a sample and not valid JSON. To use it, please:
# 1. Remove all comments. (Comments begin with #)
# 2. Change the file extension to .json.

{
  "customer_id": "network-operator",
  "design_id": "topo",
  "instance_id": "topology",
  "operation": "create",
  "topo": {
    "pop": [
      {
        "name": "e57ce4ee-9ef7-4e3c-a15d-66ab16d8f247", #this can be fetched from the
sites page, see point 1 below (Site ID)
        "pe": [
          {
            "access": [
              {
                "ce": "ce1", #this is CE device reference
                "name": "ge-0/0/3", #this is the PE interface that connects to
the CE device
                "speed": "100000",
                "type": "ethernet"
              }
            ],
            "bandwidth": "40000000000",
            "mac_addr": "1000000",
            "name": "00000000-0000-0000-1000-2c6bf5efb500", #this can be fetched
from the inventory page, see point 2 below (Device ID)
            "routes": "1000000"
          }
        ]
      }
    ]
  },
}
```

```

        "postal_code_matches": [
            {
                "name": "M1", #User defined string, has to be unique within POPs
                "regex": "94..." #this is the regex for the postal codes this POP will
accept
            }
        ]
    },
    {
        "name": "3f6ae91b-02f8-4bf3-9d4c-8cbeaf832b2c",
        "pe": [
            {
                "access": [
                    {
                        "ce": "ce2",
                        "name": "ge-0/0/3",
                        "speed": "100000",
                        "type": "ethernet"
                    },
                    {
                        "ce": "ce1",
                        "name": "ge-0/0/4",
                        "speed": "100000",
                        "type": "ethernet"
                    }
                ],
                "bandwidth": "40000000000",
                "mac_addrs": "1000000",
                "name": "00000000-0000-0000-1000-2c6bf5598200",
                "routes": "1000000"
            }
        ],
        # postal_code_matches are optional
        "postal_code_matches": [
            {
                "name": "M4",
                "regex": "91..."
            }
        ]
    }
]
}
}

```

- # 1. To get site ID, go to Inventory>Sites. If you don't see the ID column, click on the three dots above the table>Show/Hide columns and enable ID.
- # 2. To get device ID, go to Inventory>Network Inventory. If you don't see the ID column, click on the three dots above the table>Show/Hide columns and enable ID.

Sample REST API to Create VPN Resources for L3VPN, L2VPN, and L2 Circuit Services



NOTE: The operation field in the JSON file can take up the following values:

- create—Creates new network resources if none exist. However, if resources already exist, new network resources specified in the JSON file are added to the existing ones.
- modify—Overrides the existing network resources with the values passed through the JSON file.
- delete—Removes the network resources specified in the JSON file.

The following is a sample of the REST API to create a VPN resource pool for L3VPN, L2VPN, and L2 circuit services:

```
https://{server}:{port}/service-orchestration/api/v1/orgs/{ORG}/order
```

```
{
  "customer_id": "network-operator",
  "design_id": "vpn",
  "instance_id": "rds-and-rts",
  "operation": "create",
  "vpn": {
    "route_distinguisher": [
      {
        "count": 1024,
        "name": 1234
      }
    ],
    "route_target": [
      {
        "count": 1024,
        "name": 1234
      }
    ]
  }
}
```

```

        "count": 1024,
        "name": 1235
    }
],
"community": [
    {
        "name": "target:64500",
        "count": 10000
    }
],
"virtual_circuit_id": [
    {
        "start": 1,
        "count": 10000
    }
]

```

SEE ALSO

[Add Network Resource Pools and Profiles \(Day -2 Activities\) | 177](#)

[Device Onboarding Workflow | 173](#)

[Add Network Resources for Provisioning Services | 670](#)

[Service Provisioning Workflow | 648](#)

Configure Topology Resources in Network Implementation Plan

Service orchestration obtains topology resources required for configuring a service from the following options:

- Configuration in network implementation plan used to onboard the device (for example, access parameters, placement resources, and bandwidth allocated on each device, and so on).
- Configuration obtained from the element management system of the device (for example, interfaces, site ID, number of routes, and so on).
- Configuration discovered by Routing Director (for example, number of aggregated Ethernet (ae) interfaces configured on a device).



NOTE: If you do not configure the number of aggregated Ethernet interfaces on a device in a network implementation plan, Routing Director uses the number of aggregated Ethernet interfaces that it discovers on the device.

When a resource is obtained from more than one option (for example, bandwidth reserved on a device), service orchestration considers the value in the network implementation plan.

When you configure a network implementation plan with topology resource parameters, Routing Director applies the configuration (number of lag interfaces) on the device during onboarding. After the device is onboarded, you can provision EVPN, EVPN-VPWS, L3VPN, and L2 Circuit services by using the resources available on the device.

To configure topology resources in a network implementation plan:

1. A network operator or superuser must create:
 - a. A device profile for the Provider Edge (PE) devices of a service and assign the device profile to all the PE devices in that service. Note that this profile must be named as 'PE' and it must be the first profile applied to all the PE devices in the service.
 - b. In the network implementation plan, you must define the access parameters for the PE-CE interfaces (In the Configure Port page) without the IP address configuration.
 - c. After the device is onboarded, move the device to In-Service. See ["Move a Device to Production" on page 268](#).

The device can now be used as an infrastructure resource in a service order for provisioning EVPN, EVPN-VPWS, L3VPN, and L2 Circuit services.

2. Routing Director creates a Topo service order by default during device onboarding. This service order is updated with the resources (bandwidth to be allocated on each device, PE-CE access parameters, interfaces, VLANs, MAC addresses, routes, speed, and so on) assigned to the Network Implementation Plan used for onboarding the device.

The network operator or a superuser may optionally create another Topo service order with the details of the site with postal code matches information, where the device is onboarded. The postal code matches represent the postal code of the CEs connected to the site.

3. After the device is in-service, view the topology resources on the View Network Resources page (**Orchestration > Resource Instances > More > View Network Resources**). by selecting the **topo** resource instance for the **network-operator** customer and clicking the **Update** option.

When you provision a service or update a network implementation plan, Routing Director automatically updates the network resources based on the data in the network implementation plan.

RELATED DOCUMENTATION

[Configure Resource Pools for Resource Instances | 745](#)

[Add an L2 Circuit Service Instance | 810](#)

[Add an EVPN or EVPN-VPWS Service Instance | 791](#)

[Add an L3VPN Service Instance | 773](#)

Configure Resource Pools for Resource Instances

IN THIS SECTION

- [Configure a Network Resource Pool | 745](#)
- [Create a Routing Resource Pool | 746](#)
- [Create a Layer 3 Resource Pool | 748](#)
- [Create a Layer 2 Resource Pool | 750](#)
- [Create a Topology Resource Pool | 752](#)
- [Create a VPN Resource Pool | 763](#)

After you add a new resource instance, you must configure the resource pool for the instance. You can configure resource pools by using the Routing Director GUI. You can upload a preconfigured JSON file or use the GUI fields to create a resource pool.

Configure a Network Resource Pool

To configure a network resource pool for a resource instance:

1. Add a resource instance. See ["Add a Resource Instance" on page 674](#).
The Modify *Resource-Instance-Name* page appears and displays the Resource Editor and Inventory panes.
2. You can configure resource pools in either one of the following ways:
 - Upload a preconfigured JSON file that contains the resource configurations.
 - a. Click **Upload** on the Resource Editor pane.
 - b. Browse and select the preconfigured JSON file on your computer.
 - c. Click **Open** to upload the file. The values specified in the file are automatically populated in the corresponding GUI fields.

- Enter the values in the GUI fields to create the resource pool.
- a. Configure the following resource pools for the corresponding resource designs:
 - To configure a routing resource pool, see ["Create a Routing Resource Pool" on page 746](#).
 - To configure a Layer 3 resource pool, see ["Create a Layer 3 Resource Pool" on page 748](#).
 - To configure a Layer 2 resource pool, see ["Create a Layer 2 Resource Pool" on page 750](#).
 - To configure a topology resource pool, see ["Create a Topology Resource Pool" on page 752](#).
 - To configure a VPN resource pool, see ["Create a VPN Resource Pool" on page 763](#).
- b. (Optional) Click **Download** on the Resource Editor pane of the Modify *Resource-Instance-Name* page to download and save the VPN resource configuration in the JSON format on your local computer.
- 3. (Optional) Click **Download** on the Inventory pane to download and save the resource inventory in the JSON format on your local computer.
- 4. Click **Proceed**.
The Compare Resource Definition page appears.
You can view the configuration updates you made in the current resource instance in comparison to the previously saved resource instance configurations in the database.
- 5. Click **Inline** for an inline view of the comparisons and click **Side-by-side** for an adjacent view.
- 6. Click **Save and Commit**.
A service order is generated and the Monitor Order Status page appears. The page displays the execution status of the service order.
- 7. Click the *service-order-name* hyperlink to monitor the execution state of the service order. See ["View Workflow Run Details" on page 885](#).
- 8. Click the **Restore instance** hyperlink if the service order fails to commit.
An order is generated to restore the last service order for the resource instance and the Monitor Order Status page appears. The page displays the execution status of the order.



NOTE: A modify service order for a resource instance is saved but fails to commit if you modify resources that are in use for a service.

- 9. Click **OK**.

Create a Routing Resource Pool

A routing resource pool defines the Autonomous System (AS) details, BGP route reflector clusters, and Segment Identifiers (SIDs). The routing resource pool must be configured for provisioning the infrastructure service.

To create a routing resource pool using the GUI fields:

1. Add a routing resource instance. See ["Add a Resource Instance" on page 674](#).
The Modify *Resource-Instance-Name* page appears and displays the Resource Editor and Inventory panes.
2. On the Resource Editor pane of the *Resource-Instance-Name* page, expand **routing**.
3. Click the add (+) icon above the Autonomous System table.
The Autonomous System page appears.
4. Enter the values by referring to the following table:

Table 164: Fields on the Autonomous System Page

Field	Description
Name	Enter the AS number of the network in which you want to configure your infrastructure or VPN service. Enter a value in the range 0 through 4,294,967,295.
Count	Enter the number of routes available in the specified AS. Enter a value in the range 0 through 65,535.

5. Click **OK**.
The autonomous system details are listed in the Autonomous System table.
6. (Optional) To edit or delete an entry, select the entry and click the edit or delete icons present above the table.
7. Expand **Route Reflector** and click the add (+) icon above the Clusters table.
The Clusters page appears.
Use this page to manage the BGP Route Reflector Cluster IDs available in the resource pool for auto-allocation when provisioning a service.
8. In the Cluster field, enter the cluster ID that the BGP route reflector supports.
A route reflector can support multiple clusters. You can enter multiple cluster IDs for a route reflector.
9. Click **OK**.
The cluster details are listed in the Clusters table.
10. Expand **Spring > Sids**.
11. In the **Size** field, enter the number of SIDs available in the resource pool.
Enter a value in the range 0 through 65,535.
12. Expand **Base** and enter the values by referring to the following table:

Table 165: SIDs Base Parameters

Field	Description
IPv4	Enter the starting or base value for IPv4 SIDs in the routing resource pool. Enter a value in the range 0 through 65,535.
IPv6	Enter the starting or base value for IPv6 SIDs in the routing resource pool. Enter a value in the range 0 through 65,535.

13. Go to "2.b" on page 746 of the "Configure a Network Resource Pool" on page 745 procedure to save and download the resource configuration and proceed with the resource instance edition.

Create a Layer 3 Resource Pool

A Layer 3 resource pool defines resources such as IPv4 and IPv6 prefixes and loopback prefixes. The Layer 3 resource pool must be configured for provisioning infrastructure services.

To create a Layer 3 resource pool:

1. Add a Layer 3 (**L3-Addr**) resource instance. See "Add a Resource Instance" on page 674.
The Modify *Resource-Instance-Name* page appears and displays the Resource Editor and Inventory panes.
2. On the Resource Editor pane of the Modify *Resource-Instance-Name* page, expand **l3-addr**.
3. Click the add (+) icon above the IPv4 Prefixes table.
The IPv4 Prefixes page appears.
4. Enter the values by referring to the following table:

Table 166: Fields on the IPv4 Prefixes Page

Field	Description
Name	Enter a name to uniquely identify the IPv4 address pool.
Prefix	Enter the IPv4 prefix in CIDR notation (<ipv4-address>/<prefix length>). For example, 10.10.11/24.

5. Click **OK**.

The IPv4 prefix details are listed in the IPv4 Prefixes table.

6. Click the add (+) icon above the IPv6 Loopbacks table.
The IPv6 Loopbacks page appears.

7. Enter the values by referring to the following table:

Table 167: Fields on the IPv6 Loopbacks Page

Field	Description
Name	Enter a name to uniquely identify the loopback interface.
Prefix	Enter the loopback prefix (in CIDR notation) to define the range of IPv6 addresses that can be assigned to the loopback interface.

8. Click **OK**.
The IPv6 loopback details are listed in the IPv6 Loopbacks table.

9. Click the add (+) icon above the IPv6 Prefixes table.
The IPv6 Prefixes page appears.

10. Enter the values by referring to the following table:

Table 168: Fields on the IPv6 Prefixes Page

Field	Description
Name	Enter a name to uniquely identify the IPv6 address pool.
Prefix	Enter the IPv6 prefix in CIDR notation (<ipv6-address>/<prefix length>).

11. Click **OK**.
The IPv6 prefix details are listed in the IPv6 Prefixes table.

12. Click the add (+) icon above the Loopbacks table.
The Loopbacks page appears.

13. Enter the values by referring to the following table:

Table 169: Fields on the Loopbacks Page

Field	Description
Name	Enter a name to uniquely identify the loopback interface.
Prefix	Enter the loopback prefix (in CIDR notation) to define the range of IPv4 addresses that can be assigned to the loopback interface.

14. Click **OK**.
The loopback details are listed in the Loopbacks table.
15. (Optional) To edit or delete an entry in any of the tables, select the entry and click the edit or delete icons present above the table.
16. Go to ["2.b" on page 746](#) of the ["Configure a Network Resource Pool" on page 745](#) procedure to save and download the resource configuration and proceed with the resource instance edition.

Create a Layer 2 Resource Pool

A Layer 2 resource pool defines resources such as Ethernet VPN (EVPN) Ethernet Segment Identifiers (ESI), LACP administrative keys, and LACP system identifiers. The Layer 2 resource pool must be configured for provisioning EVPN services.

To create a Layer 2 resource pool:

1. Add a Layer 2 (**L2-Addr**) resource instance. See ["Add a Resource Instance" on page 674](#).
The Modify *Resource-Instance-Name* page appears and displays the Resource Editor and Inventory panes.
2. On the Resource Editor pane of the Modify *Resource-Instance-Name* page, expand **l2-addr**.
3. Enter the EVPN ESI details:
 - a. Click the add (+) icon above the EVPN ESI table.
The EVPN ESI page appears.
 - b. Enter the values by referring to the following table:

Table 170: Fields on the EVPN ESI Page

Field	Description
Name	Enter a unique Ethernet segment identifier (ESI). An Ethernet segment refers to the set of Ethernet links that connect a multihomed CE device to multiple PE devices. The ESI uniquely identifies an Ethernet segment in the network.
Count	Enter the number of EVPN ESIs that are available in the resource pool.

- c. Click **OK**.

The EVPN ESI details are listed in the EVPN ESI table.

4. Enter LACP administrative key details:

- a. Expand **LACP** and click the add (+) icon above the Admin Key table.

The Admin Key page appears.

- b. Enter the values by referring to the following table:

Table 171: Fields on the Admin Key Page

Field	Description
Start	Enter the starting value for the number of LACP administrative keys available in the resource pool. Enter a value in the range 0 through 65,535. You must specify LACP administrative keys for Multichassis Link Aggregation (MC-LAG) or Link Aggregation Groups (LAGs) that span physical ports across two separate devices.
Size	Enter the number of LACP administrative keys available for auto-allocation in the resource pool, beginning from the start value. Enter a value in the range 1 through 65,535.

- c. Click **OK**.

The LACP administrative key details are listed in the Admin Key table.

5. Enter the LACP system ID details:
 - a. Expand **LACP** and click the add (+) icon above the System ID table.
The System ID page appears.
 - b. Enter the values by referring to the following table:

Table 172: Fields on the System ID Page

Field	Description
Name	Enter the unique LACP system ID that is available for auto-allocation in the L2 resource pool. You must specify LACP system ID for Multichassis Link Aggregation (MC-LAG) or LAGs that span physical ports across two separate devices.
Count	Enter the number of ports to which the unique LACP system ID is assigned. Enter a value in the range 1 through 255.

- c. Click **OK**.
The LACP system ID details are listed in the System ID table.
6. (Optional) To edit or delete an entry in any of the tables, select the entry and click the edit or delete icons present above the table.
7. Go to ["2.b" on page 746](#) of the ["Configure a Network Resource Pool" on page 745](#) procedure to save and download the resource configuration and proceed with the resource instance edition.

Create a Topology Resource Pool

A topology resource pool defines resources such as the sites, and PE device parameters such as access, bandwidth, routes, and postal code matches. The topology resource pool must be configured for provisioning L3VPN, EVPN, EVPN-VPWS, and L2 circuit services.



NOTE: If you configure topology resources by specifying them in the network implementation plan, you must not add topology resource pools manually. See ["Configure Topology Resources in Network Implementation Plan" on page 743](#).
The default topology resource instance managed through the network implementation plan is a read-only instance and cannot be modified.

To create a topology resource pool:

1. Add a topology resource instance. See ["Add a Resource Instance" on page 674](#).
The Modify *Resource-Instance-Name* page appears and displays the Resource Editor and Inventory panes.
2. On the Resource Editor pane of the Modify *Resource-Instance-Name* page, expand **topo**.
3. Click the add (+) icon above the POP/Site table.
The POP page appears.
A point-of-presence (POP) refers to a site or physical location that hosts the PE devices.
4. Click the **Name** drop-down list and select the site from the list of available sites.
The drop-down lists all the sites that you add by using the Sites page (**Inventory > Common Resources > Sites**). See ["About the Sites Page" on page 88](#).
5. Enter the node details. Click the add (+) icon above the Node table.
The Node page appears.
 - a. Select the PE device name, and configure supported bandwidth, MAC addresses, roles, routes, and IRB details by referring to the following table:

Table 173: Fields on the Node Page

Field	Description
Name	Click the Name drop-down and select the PE device from the list of available devices.
Bandwidth	Enter the aggregated bandwidth capacity (in Mbps) of the PE device. Aggregated bandwidth is the combined bandwidth capacity of all the interfaces on the device.
Max MAC addresses	Enter the number of MAC addresses that the PE device can manage. The MAC address count is used while allocating topology resources for EVPN and EVPN-VPWS service provisioning.
Roles	Click the Roles drop-down and select the node resource role from the list of available nodes. PE is the only node resource role supported in this release.

Table 173: Fields on the Node Page (Continued)

Field	Description
Max Routes	Enter the number of routes that the VRF table of the PE device can manage. The route capacity is used while allocating topology resources for L3VPN service provisioning.
Expand IRBs	
Start	Enter the starting IRB for the range of IRB IDs to which the PE interface connects. Enter a value in the range 1 through 65535.
Size	Enter the number or range of IRB IDs available for a PE device. You can define a range using a starting ID and the number of IRBs required. Enter a value in the range 1 through 65535.

- b. Click the add (+) icon above the Access Interfaces table.
The Access page appears.
- c. Configure the PE interface name, type and speed, and specify whether the PE interface connects to a single or multiple CE devices, by referring to the following table:

Table 174: Fields on the Access Page

Field	Description
Name	Enter the name of the physical PE interface that connects to the CE interface. For example, et-0/0/0. If you are configuring multihoming on EVPN VPWS services, then enter the pre-provisioned link aggregation group (LAG) Aggregated Ethernet interface name. For example, ae1. NOTE: To configure multihoming you must have pre-provisioned the LAG interface on the device. You can pre-provision a LAG interface through the Routing Director GUI by using network implementation plans. Alternatively, you can log in to the device and manually configure the LAG interface and its member links at the IFD level on the device. For example: - For tagged interfaces: <pre> set interfaces ae1 flexible-vlan-tagging set interfaces ae1 encapsulation flexible-ethernet-services set interfaces ae1 aggregated-ether-options minimum-links 1 set interfaces ae1 aggregated-ether-options lacp system-id 00:00:00:00:11:02 set interfaces ae1 aggregated-ether-options lacp admin-key 10001 set interfaces et-0/0/2 gigether-options 802.3ad ae1 </pre> - For untagged interfaces: <pre> set interfaces ae1 encapsulation ethernet-ccc set interfaces ae1 aggregated-ether-options minimum-links 1 set interfaces ae1 aggregated-ether-options lacp system-id 00:00:00:00:11:02 set interfaces ae1 aggregated-ether-options lacp admin-key 10001 set interfaces et-0/0/2 gigether-options 802.3ad ae1 </pre>

Table 174: Fields on the Access Page (*Continued*)

Field	Description
CE Spec	Specify if you want to connect the PE interface to a single CE device or multiple CE devices. Click the drop-down and select one of the following options: • ce—The PE interface connects to a single CE device. When you select ce, the CE and Tag Type fields appear. You can enter the CE name and select a tag type by referring to step 5.d. • ces—The PE interface connects to multiple CE devices. When you select ces, the Ces section appears. You can enter the CE names and select a tag type by referring to step 5.e. NOTE: • Enter VLAN details if you want the access interface type in tagged mode. A tagged interface receives and sends Ethernet frames with VLAN tags. • If you optionally specify a VLAN while creating a service instance, service orchestration automatically selects an access interface on a PE device that has the specified VLAN available for use.
Lag	Toggle on the Lag button to True to indicate that the interface is an Aggregated Ethernet interface. Enable this button when you configure multihoming in EVPN VPWS services. When the Lag toggle button is enabled, service orchestration only allocates the pre-provisioned LAG interfaces for multihoming site network access. Note that, you must have pre-provisioned the LAG interface on the device. If the Lag option is enabled and you configure a single-homed EVPN VPWS service, the Aggregated Ethernet interface behaves as a normal physical interface and can be allocated for single-homed site network access, as well.

Table 174: Fields on the Access Page (*Continued*)

Field	Description
Speed	Enter the speed (in Mbps) of the physical PE interface. The interface speed refers to the maximum data transmission rate that the interface supports. The value entered here is displayed under the Placement option once placement options are updated in the EVPN VPWS service instance. The value entered here does not affect the topology resource.
Type	Select the PE access type from the drop-down. Access type defines how the PE interface connects to the CE device. Ethernet is the only access type that is supported in this release.

d. (For single CE connection) Enter the VLAN details for single CE device connection:

- i. Select **ce** from the **CE Spec** drop-down.
- ii. (Optional) Enter a CE device name to uniquely identify the device in the network.
- iii. Select a tag type from the **Tag Type** drop-down. You can select **dot1q**, **qinq**, or **untagged** as the tag type.

The default tag type is **untagged**.

1. When you select **dot1q** from the **Tag Type** drop-down, the VLANs table appears. Click the add (+) icon above the VLANs table, the VLANs page appears.

Table 175: Dot1q VLAN Details for Single CE Connection on the VLANs Page

Field	Description
Size	Enter the range of VLANs from the CE device that the PE interface supports. Enter a value in the range 1 through 4094.

Table 175: Dot1q VLAN Details for Single CE Connection on the VLANs Page
(Continued)

Field	Description
Start	Enter the starting VLAN ID for the range of VLANs from the CE device that the PE interface supports. Enter a value in the range 1 through 4094. For example, if 200 is the starting VLAN ID, the first VLAN supported by the PE interface has a VLAN ID of 200.



NOTE: To configure multihoming in EVPN VPWS services with pre-provisioned LAG interfaces, select **dot1q** or **untagged**.

Click **OK**. The VLAN details are listed in the VLANs table.

- When you select **qinq** from the **Tag Type** drop-down, the Inner VLANs and Outer VLANs tables appear.
 - Click the add (+) icon above the Inner VLANs table. The Inner VLANs page appears.

Table 176: Inner VLAN Details for Single CE Connection on the Inner VLANs Page

Field	Description
Start	Enter the starting VLAN ID for the range of inner VLANs. Enter a value in the range 1 through 4094. For example, if 200 is the starting inner VLAN ID, the first inner VLAN supported by PE device has a VLAN ID of 200.
Size	Enter the range for inner VLANs. Enter a value in the range 1 through 4094.

Click **OK**. The details are listed in the Inner VLANs table.

- Click the add (+) icon above the Outer VLANs table. The Outer VLANs page appears.

Table 177: Outer VLAN Details for Single CE Connection on the Outer VLANs Page

Field	Description
Start	Enter the starting VLAN ID for the range of outer VLANs. Enter a value in the range 1 through 4094. For example, if 200 is the starting outer VLAN ID, the first outer VLAN supported by PE device has a VLAN ID of 200.
Size	Enter the range for outer VLANs. Enter a value in the range 1 through 4094.

Click **OK**. The details are listed in the Outer VLANs table.

- Select **untagged** from the **Tag Type** drop-down to assign no VLANs to the interface. An untagged interface receives and sends Ethernet frames without VLAN tags.

If you do not select any tag type, the interface will display only untagged interfaces.

- (For multiple CE connections) Enter VLAN details for multiple CE connections.

You must add the device name, select a tag type, and enter VLAN details for all the CE devices that connect to the PE interface.

- Select **ces** from the **CE Spec** drop-down.
- Click the add (+) icon above the CEs table.

The CEs page appears.

- (Mandatory) Enter CE device names to uniquely identify the devices in the network.
- Select a tag type from the **Tag Type** drop-down. You can select **dot1q**, **qinq**, or **untagged** as the tag type.

The default tag type is **untagged**.

- When you select **dot1q** from the **Tag Type** drop-down, the VLANs section appears.

Enter VLANs details by referring to the following table:

Table 178: Dot1q VLAN Details for Multiple CE Connections on the CEs Page

Field	Description
Size	Enter the range of VLANs from the CE device that the PE interface supports. Enter a value in the range 1 through 4094.
Start	Enter the starting VLAN ID for the range of VLANs from the CE device that the PE interface supports. Enter a value in the range 1 through 4094. For example, if 200 is the starting VLAN ID, the first VLAN connected from the CE device to the PE interface has a VLAN ID of 200.



NOTE: To configure multihoming in EVPN VPWS services with pre-provisioned LAG interfaces, select **dot1q** or **untagged**.

Click **OK**. The device names and VLAN details are listed in the CEs table.

- When you select **qinq** from the **Tag Type** drop-down, the Inner VLANs and Outer VLANs tables appear. Enter Inner VLANs and Outer VLANs details by referring to the following table:

Table 179: Q-in-Q VLAN Details for Multiple CE Connections on the CEs Page

Field	Description
Inner VLANs	
Size	Enter the range for inner VLANs. Enter a value in the range 1 through 4094.

Table 179: Q-in-Q VLAN Details for Multiple CE Connections on the CEs Page
(Continued)

Field	Description
Start	Enter the starting VLAN ID for the range of inner VLANs. Enter a value in the range 1 through 4094. For example, if 200 is the starting inner VLAN ID, the first inner VLAN supported by the PE interface has a VLAN ID of 200.
Outer VLANs	
Size	Enter the range for outer VLANs. Enter a value in the range 1 through 4094.
Start	Enter the starting VLAN ID for the range of outer VLANs. Enter a value in the range 1 through 4094. For example, if 200 is the starting outer VLAN ID, the first outer VLAN supported by the PE interface has a VLAN ID of 200.

Click **OK**. The Inner VLANs and Outer VLANs details are listed in the CEs table.

3. Select **untagged** from the **Tag Type** drop-down to assign no VLANs to the interface. An untagged interface receives and sends Ethernet frames without VLAN tags.

If you do not select any tag type, the interface will display only untagged interfaces.

- v. (Optional) For a quick JSON view of the VLAN details, hover over the **Show details** hyperlink for a VLAN.
- f. Click **OK** on the Access page.
The access details are listed in the Access Interfaces table.
- g. Click the add (+) icon above the Aggregated Ethernet Interfaces table.
The AE page appears.

- h. Enter the AE interface name and click **OK**.

The AE interface name is listed in the AE table.

You can add multiple interface names for all the AE interfaces that are available on the PE device to provision services.



NOTE: If the AE interface count for the device is not configured in the network implementation plan, you can configure the AE interfaces count on the device by executing the following command on the device CLI:

```
root@device-name# set chassis aggregated-devices ethernet device-count <device-count>
```

- i. Click **OK** on the Node page.

The node details are listed in the Node table.



NOTE: PE section is deprecated.

6. Enter the postal code match details:

- a. Click the add (+) icon above the Postal Code Matches table.

The Postal Code Matches page appears.

- b. Enter the values by referring to the following table:

Table 180: Fields on the Postal Code Matches Page

Field	Description
Name	Enter the name of the site where the PE device is located.
Country Code	Enter the ISO alpha-2 code of the country where the site is located. For example, ZA for South Africa and CH for Switzerland.
Regex	Enter the postal code regex of the site where the PE device is located.

- c. Click **OK** on the Postal Code Matches page.

The postal code details are listed in the Postal Code Matches table.

7. (Optional) To edit or delete an entry in any of the tables, select the entry and click the edit or delete icons present above the table.

8. Click **OK** on the POP page.

The site or POP name, PE details, and postal code match details are listed in the POP/Site table.

To view PE and postal code details for a POP in the JSON format, hover over the PE and postal code match hyperlinks.

9. Go to ["2.b" on page 746](#) of the ["Configure a Network Resource Pool" on page 745](#) procedure to save and download the resource configuration and proceed with the resource instance edition.



NOTE: PE nodes and interfaces that connect to CE devices cannot be configured simultaneously in the network implementation plan and in the topology resource pool for provisioning infrastructure services. You must delete PE nodes and interfaces parameters from the topology resource pool if you are defining the same parameters in the network implementation plan.

To delete the PE nodes and interfaces parameters from the topology resource pool, modify the topology resource instance and upload the modified instance to the Juniper Routing Director database. See ["Modify a Resource Instance" on page 766](#).

Create a VPN Resource Pool

A VPN resource pool defines resources such as the route distinguishers, route targets and route target community details, and virtual circuits. The VPN resource pool must be configured for provisioning L3VPN, EVPN, or L2 circuit services.

To add a VPN network resource pool:

1. Add a VPN resource instance. See ["Add a Resource Instance" on page 674](#).
The Modify *Resource-Instance-Name* page appears and displays the Resource Editor and Inventory panes.
2. On the Resource Editor pane of the Modify *Resource-Instance-Name* page, expand **vpn-resources**.
3. Click the add (+) icon above the Community table.
The Community page appears.
4. Enter the values by referring to the following table:

Table 181: Fields on the Community Page

Field	Description
Name	Enter a unique name for the route target community.

Table 181: Fields on the Community Page (Continued)

Field	Description
Count	Enter the number of route targets included in the target community. Enter a value in the range 1 through 65,535.

5. Click **OK**.
The route target community details are listed in the Community table.
6. Click the add (+) icon above the Route Distinguisher table.
The Route Distinguisher page appears.
7. Enter the values by referring to the following table:

Table 182: Fields on the Route Distinguisher Page

Field	Description
Name	Enter a unique ID for the route distinguisher. Enter a value in the range 0 through 65,535.
Count	Enter the number of route distinguishers with the unique ID available in the resource pool. Enter a value in the range 1 through 65,535.

8. Click **OK**.
The route distinguisher details are listed in the Route Distinguisher table.



NOTE: Route distinguishers are automatically allocated if you do not define the route distinguisher details in the VPN resource pool.

9. Click the add (+) icon above the Route Target table.
The Route Target page appears.
10. Enter the values by referring to the following table:

Table 183: Fields on the Route Target Page

Field	Description
Name	Enter a unique ID for the route target. Enter a value in the range 0 through 65,535.
Count	Enter the number of route targets with the unique ID available in the resource pool. Enter a value in the range 1 through 65,535.

11. Click **OK**.
The route target details are listed in the Route Target table.
12. Click the add (+) icon above the Virtual Circuit ID table.
The Virtual Circuit ID page appears.
13. Enter the values by referring to the following table:

Table 184: Fields on the Virtual Circuit ID Page

Field	Description
Start	Enter the starting value for the range of virtual circuit IDs available in the resource pool. Enter a value in the range 0 through 4,294,967,295.
Count	Enter the number of virtual circuit IDs available in the resource pool. Enter a value in the range 1 through 4,294,967,295.

14. Click **OK**.
The virtual circuit details are listed in the Virtual Circuit ID table.
15. (Optional) To edit or delete an entry in any of the tables, select the entry and click the edit or delete icons present above the table.
16. Go to ["2.b" on page 746](#) of the ["Configure a Network Resource Pool" on page 745](#) procedure to save and download the resource configuration and proceed with the resource instance edition.

SEE ALSO

[Resource Instance Overview](#) | 669

[Add a Resource Instance](#) | 674

[Modify, Delete, and Restore Resource Instances](#) | 766

Modify, Delete, and Restore Resource Instances

IN THIS SECTION

- [Modify a Resource Instance](#) | 766
- [Delete a Resource Instance](#) | 767
- [Restore a Resource Instance](#) | 767

Use the following procedures to modify, delete, and restore resource instances from the Routing Director database.

Modify a Resource Instance

To modify a resource instance:

1. Click **Orchestration > Resource Instances**.
The Resource Instances page appears.
2. In the Resource Instances table, select the instance you want to modify and click the edit (pencil) icon.
The Modify *Resource-Instance-Name* page appears.
3. Modify the resource instance by uploading a preconfigured JSON file with the updates, or by editing the GUI fields in the Resource Editor pane. See "[Configure a Network Resource Pool](#)" on page 745.
After the modify service order for the resource instance is successfully executed, the updated configurations are uploaded to the Routing Director database.



NOTE: If you modify a resource instance when the resources configured in the instance are in use for a service, the modify service order is saved but fails to commit. When the service order commit fails, you can restore the last service order for the resource instance. See "[Restore a Resource Instance](#)" on page 767.

Delete a Resource Instance

To delete a resource instance:

1. Click **Orchestration > Resource Instances**.

The Resource Instances page appears.

2. In the Resource Instances table, select the instance you want to delete and click the edit (pencil) icon.

3. Click **More > Deprovision**.

Alternatively, right-click the resource instance name and click **Deprovision**.

4. On the **Deprovision Resource Instance** dialog box, click **Yes** to confirm.

You see a message that the resource instance is successfully deleted.



NOTE: If you delete a resource instance when the resources configured in the instance are in use for a service, the delete service order fails to execute. When the delete order fails, you can restore the last service order for the resource instance. See ["Restore a Resource Instance" on page 767](#).

Restore a Resource Instance

If you modify or delete a resource instance when the resources configured in the instance are in use for a service, the corresponding modify or delete service order that is generated is saved but fails to commit. When the service-order commit fails, you can restore the last service order for the resource instance.

To restore the last service order for a resource instance:

1. Click **Orchestration > Resource Instances**.

The Resource Instances page appears.

2. In the Resource Instances table, select the instance you want to restore.

3. Click **More > Restore**.

The **Restore** option is enabled only when a resource instance fails to commit. This option is disabled if the resource instance is successfully committed.

4. Click **Yes** on the Restore Resource Instance dialog box to confirm.

A service order is generated and the Monitor Order Status page appears. The page displays the execution state of the order.

Alternatively, you can restore the last service order for a resource instance by clicking the **Restore instance** hyperlink on the Monitor Order Status page when a service order fails to commit after you modify the network resource pool. For more information, see step 8 of the ["Configure a Network Resource Pool" on page 745](#) procedure.

SEE ALSO

[About the Resource Instances Page | 672](#)

[Add a Resource Instance | 674](#)

[Configure Resource Pools for Resource Instances | 745](#)

Manage Placement Configurations for Service Instances

IN THIS SECTION

- [Manage Placement Configurations by Using the Routing Director GUI | 768](#)
- [Manage Placement Configurations by Using the Service Orchestration cMGD CLI | 770](#)

Before provisioning a service in the network, you (superuser or network administrator) must configure and upload network resource pools to the Routing Director database. You can upload network resources by using the GUI, service orchestration cMGD CLI, or by using REST API. For more information, see ["Add Network Resources for Provisioning Services" on page 670](#).

Placement refers to the allocation of network resources to provision a service. You can manage allocation of network resources such as PE devices, interfaces, CE devices, and VLANs when provisioning an EVPN, EVPN-VPWS, L3VPN, or L2 circuit service for a customer. Automatically assign all possible placement options available for a customer from the network resource pools that you configured and uploaded. After placement options are assigned, you can select and allocate devices and interfaces from the available options to provision services. You can manage allocation of network resources when creating or modifying a service instance by using the Routing Director GUI or the service orchestration cMGD CLI.

Manage Placement Configurations by Using the Routing Director GUI

Prerequisite

["Add Network Resources for Provisioning Services" on page 670](#).

You can assign placement options for an L3VPN, EVPN, EVPN-VPWS, or L2 circuit service instance when you are creating or modifying an instance by using the Routing Director GUI. Use the **Update Placements** option to automatically assign possible placement options for a service instance. Before auto allocation of resources, the service instance is saved and uploaded to the Routing Director database. After placement options are successfully assigned, you see a confirmation message on the

GUI. You can then edit the saved service instance and update placement options in the Modify *service-instance-type* page.

To assign placement options for an L3VPN, EVPN, EVPN-VPWS or L2 circuit service instance:

1. Click *Orchestration* > *Instances*.

The Service Instances page appears.

2. Assign placement options when you create a service instance:

- To add a new L3VPN service instance, see ["Add an L3VPN Service Instance" on page 773](#).

Follow step 13 of the ["Add L3VPN Service Site Details" on page 777](#) procedure to assign placement options and configure placement parameters for the service instance.

- To add a new EVPN or EVPN-VPWS service instance, see ["Add an EVPN or EVPN-VPWS Service Instance" on page 791](#).

Follow step 12 of the ["Add EVPN or EVPN-VPWS Service Site Details" on page 796](#) procedure to assign placement options and configure placement parameters for the service instance.

- To add a new L2 circuit service instance, see ["Add an L2 Circuit Service Instance" on page 810](#).

Follow step 11 of the ["Add L2 Circuit VPN Nodes" on page 814](#) procedure to assign placement options and configure placement parameters for the service instance.

3. Assign placement options when you are modifying a service instance:

- Assign placement options from the Service Instances page:

- a. Select the service instance for which you want to assign placement options and click **More** > **Update Placements**.

You see a confirmation message on the GUI after placement options are successfully assigned.

- b. Select the service instance and click the edit icon.

The Modify *service-instance-type* page appears.

- c. Proceed to the Site Settings section and follow Step 5.h of the ["Modify a Service Instance" on page 822](#) procedure to modify placement parameters for the L3VPN, EVPN, and EVPN-VPWS service instances. Follow Step 6 of the ["Modify a Service Instance" on page 822](#) procedure to modify placement parameters for the L2 circuit service instance.

- (For L3VPN, EVPN, and EVPN-VPWS) Assign placement options from the Site Settings section of the Modify *service-instance-type* page:

- a. On the Service Instance page, select the service instance for which you want to assign placement options and click the edit icon.

The Modify *service-instance-type* page appears.

- b. Proceed to the Site Settings section.
- c. Assign placement options and modify placement parameters for the service instance by following step 5.h of the ["Modify a Service Instance" on page 822](#) procedure.
- (For L2 circuit) Assign placement options from the VPN Nodes section of the Modify *service-instance-type* page:
 - a. On the Service Instance page, select the service instance for which you want to assign placement options and click the edit icon.

The Modify *service-instance-type* page appears.

- b. Proceed to the VPN Nodes section.
- c. Assign placement options and modify placement parameters for the service instance by following Step 6 of the ["Modify a Service Instance" on page 822](#) procedure.

Manage Placement Configurations by Using the Service Orchestration cMGD CLI

Superusers with root user privileges can use the service orchestration cMGD CLI to automatically assign placement options for a service instance.

To configure placement options for a service instance:

1. Log in to the service orchestration cMGD CLI. See ["Access the Service Orchestration cMGD CLI" on page 1577](#).
2. Execute the request service order place *service-order-name* modifier options command to assign possible placement configurations for the service.
The service order status changes to **optioned**.
Modifier refers to placement modifier that changes placement configurations for a service.
3. Execute the request service order sync *service-order-name* command to load and commit the configuration updates. See ["request service order sync" on page 1630](#).
4. Execute the show configuration foghorn:customers customer *customer-id* command to see the placement configuration options for the specified customer. For example, all possible site location options. See ["show configuration foghorn:customers" on page 1610](#).
Repeat steps 2 to 4 to see more placement options for the customer, such as ports, interfaces, and VLANs.
5. Execute the request service order sync *service-order-name* command to load and commit the configuration updates.
6. Execute the request service order place *service-order-name* modifier place command. This command selects placement options and creates the placement configurations for the service.
To see the service order status, execute the show service order status command. The service order status changes to **placed**.

To modify placement parameters for a service instance, select and allocate placement parameters from the options that were automatically assigned by using the `request service order place` command.

SEE ALSO

[Add an L3VPN Service Instance | 773](#)

[Add an EVPN or EVPN-VPWS Service Instance | 791](#)

[Add an L2 Circuit Service Instance | 810](#)

[Modify a Service Instance | 822](#)

[request service order place | 1634](#)

[request service order sync | 1630](#)

[show configuration foghorn:customers | 1610](#)

[show service order status | 1591](#)

View Network Resources

Network resources are device parameters that identify the device in a network. For example, IP addresses, autonomous system number, BGP cluster ID, and segment identifiers (SID). For a device, you can configure Routing Director to automatically assign the values for the network resources during onboarding. For Routing Director to assign values for the resources automatically, you must upload values for the network resources to Routing Director. See ["Add Network Resource Pools" on page 676](#).

Routing Director enables you to view the network resources that are available for assignment and that are used up.

To view network resources:

1. Log in to Routing Director.
The Select an organization page appears.
2. Click the organization in which you want to view the network resources.
The Troubleshoot Devices page appears.
3. Navigate to **Orchestration > Service > Resource Instances**.
The Resource Instances page appears.
4. Click **More > View Network Resources**.
The Network Resources page appears. On this page, you can view:
 - Resource pools defined for the different resources
 - Number of values used up in each resource pool

- Number of values that are available for use in each resource pool

RELATED DOCUMENTATION

[Device Onboarding Workflow | 173](#)

[Network Implementation Plan Overview | 230](#)

Provision VPN Services

IN THIS CHAPTER

- Add an L3VPN Service Instance | 773
- Add L3VPN Service Site Details | 777
- Add an EVPN or EVPN-VPWS Service Instance | 791
- Add EVPN or EVPN-VPWS Service Site Details | 796
- Add an L2 Circuit Service Instance | 810
- Add L2 Circuit VPN Nodes | 814
- Modify a Service Instance | 822

Add an L3VPN Service Instance

A superuser or network administrator can use Routing Director to provision an L3VPN service in their network.

When you create and save an L3VPN service instance, Routing Director generates a create service order. After you provision the service instance, Routing Director activates the automated workflow for fulfilling the service order and provisions the service in the network.

You can create an L3VPN service instance by uploading a preconfigured JSON file or by entering the details in the UI fields on the Add L3VPN Service page.

To create an L3VPN service instance:

1. Click **Orchestration > Instances**.
The Service Instances page appears.
2. Click **Add > L3VPN**.
The Add L3VPN Service page appears.
3. In the General section, enter the values by referring to [Table 185 on page 775](#).
4. Click **Next** to proceed to the **Site Settings** section.

5. Enter site-specific details in the Site Settings section. See ["Add L3VPN Service Site Details" on page 777](#).



NOTE: If you upload a preconfigured file, the values specified in the file are automatically populated in the corresponding UI fields.

6. Click **Next** to proceed to the **Summary** section.

You see a graphical map of your service topology and a summary of the instance details in the **Summary** section.

7. (Optional) Click **Export** to export and save the service details in JSON format on your local system.
8. (Optional) Click **Reset Placements** to remove existing placement configurations for the service instance.

Provisioning of a service instance might fail when existing placement configurations for the service instance are no longer available or applicable. When you reset placement parameters, service orchestration automatically assigns placement options for the instance before provisioning the service.

9. Click **Next** to proceed to the Schedule section.

10. (Optional) Switch the Schedule toggle on if you want to schedule when to provision the service.

The Schedule option is disabled by default. When you schedule provisioning, the **Save & Provision** button on top of the page is disabled and grayed out.

The Schedule field appears. Perform the following steps to schedule the date and time to provision the service:

- a. Click the calendar icon in the Schedule field.
- b. Select the date and time for provisioning the service, and click **OK**.
The scheduled date and time are displayed in the MM/DD/YYYY HH:MM AM/PM format.



NOTE:

- You must set a present or future date and time to schedule provisioning of a service. If you select a past time and save the service order, the service is provisioned immediately.
- To delete a provisioning schedule, select the *service-instance-name* on the Service Instances page and click the edit icon. On the Modify *Service-Instance-Name* page, proceed to the Schedule section. In the Schedule section, switch the Schedule toggle off and click **Save**. The provisioning schedule is deleted and the service instance is provisioned immediately.

- Scheduling of service provisioning is a Beta feature in this release.

11. To save and provision the service instance, click one of the following buttons:

- **Save & Provision**—Click this button and confirm when prompted, to save the service instance to the Routing Director database and provision the instance immediately. This option is disabled if you have scheduled the provisioning of a service instance.
- **Save**—Click this button to save the service instance to the Routing Director database.

The Service Instances page appears and the new service instance is listed in the Service Instances table. To provision the service instance, select the service instance name and click **Provision**. Alternatively, you can schedule provisioning of the service instance as defined in step 10.



NOTE:

- The **Provision** button is disabled if you schedule the provisioning of a service instance. The service is provisioned automatically at the scheduled date and time.



NOTE: Fields marked with an asterisk (*) are mandatory.

Table 185: Fields in the Add L3VPN Service Page General Section

Field	Description
Upload JSON File	Click Browse to upload a preconfigured JSON file. You see a message that the file is successfully imported. The values specified in the file are automatically populated in the corresponding UI fields.

Table 185: Fields in the Add L3VPN Service Page General Section (*Continued*)

Field	Description
Customer*	Enter the name of the customer for whom you are provisioning the service: - If you already added the customer name by using the Customer Inventory page (Orchestration > Service > Customers), click the Customer field to see a drop-down with the customer names that you added. Select the customer name from the drop-down list. - Alternatively, click the Add Customer link to create a new customer. See "Add a Customer" on page 666. The customer name must be unique within an organization. For example, network-operator.
Instance name*	Enter a name for the service instance. For example, l3vpn. The instance name can be a set of alphanumeric characters and the special character hyphen (-). The maximum number of characters allowed is 64.
VPN Id*	Enter the ID you want to assign to the VPN. For example, vpn10. The VPN ID must not exceed 64 characters.
VPN Service Topology*	Select a topology for the VPN service: - Any-to-any topology—In this service topology, all VPN sites can exchange network traffic with each other without any restrictions. - Hub-spoke topology—In this service topology, spoke sites can exchange network traffic with hub sites, but not with other spoke sites. The hub sites can exchange network traffic with other hub sites. NOTE: The minimum release required to configure L3VPN hub and spoke topology on ACX Series routers is 24.2R1.

RELATED DOCUMENTATION

Service Provisioning Workflow | [648](#)

Add L3VPN Service Site Details | [777](#)

Add L3VPN Service Site Details

IN THIS SECTION

- [Add L3VPN Site](#) | 777
- [Add L3VPN Site Network Access Parameters](#) | 779
- [Add Access Diversity Parameters for L3VPN Site Network Access](#) | 786
- [Add Routing Protocol for L3VPN Site Network Access](#) | 788

An L3VPN service is provisioned on a collection of sites to exchange network traffic over a shared IP infrastructure. Use the Customer Site Settings section of the Add L3VPN Service page to enter details of sites that you want to connect to the VPN.

You configure the following in the Customer Site Settings section:

- General details (site ID, location, devices, and maximum routes) about the site. See "[Add L3VPN Site](#)" on page 777.
- Parameters for site network access (connection with the service provider network). See "[Add L3VPN Site Network Access Parameters](#)" on page 779.

Add L3VPN Site

Prerequisite

Before adding the L3VPN service site details, you must create a new L3VPN service instance and enter general details about the service. Follow steps 1 to 4 of the "[Add an L3VPN Service Instance](#)" on page 773 procedure to create a new instance and add L3VPN service general details.

An L3VPN site refers to the geographical area where the VPN service spans. To add an L3VPN service site in the Customer Site Settings section:

1. Click the add (+) icon on the top-right corner of the Customer Sites section.
The Add Site page appears.
2. Enter site ID, location, and device values by referring to [Table 186 on page 778](#).
3. Enter maximum routes value:

- a. Expand **Maximum Routes** and click the add (+) icon on the top-right corner of the Address Family section.

The Address Family page appears.

- b. Click the **Af** drop-down to select either IPv4 or IPv6 address family to configure routing for the virtual and forwarding (VRF) table.

- c. Enter the maximum routes value.

Enter a value in the range 0 through 4,294,967,295.

This value indicates the number of routes that the VRF table can accept for the IPv4 and IPv6 address families.

- d. Click **OK**.

The details you entered are listed in the Address Family table.

- e. (Optional) To edit or delete an entry, use the edit or delete options present above the Address Family table.

4. Do any of the following:

- Click **Cancel** to exit the Add Site page without saving the changes you made.
- Click **OK** to save the site details you added. The site you added is listed in the **Sites** table.

You can view the configured properties and the number of maximum routes by expanding **Properties**.



NOTE: Fields marked with an asterisk (*) are mandatory.

Table 186: Fields on the Add Site Page (L3VPN)

Field	Description
Site ID*	Enter a unique site ID to identify the site in the network. For example, site1. The site ID can contain alphanumeric characters and special characters hyphens (-) and underscores (_). The maximum number of characters allowed is 64.
Locations	
Location ID*	Enter a unique ID for a location in the site. For example, MAN for Manhattan and BRO for Brooklyn in the New York site.

Table 186: Fields on the Add Site Page (L3VPN) (Continued)

Field	Description
Address	Enter the address (number and street) of the location in a site. For example, 123, Avenue Street.
City	Enter the city where the site is located. For example, MAN for Manhattan.
Country Code	Enter the ISO alpha-2 code of the country where the site is located. For example, ZA for South Africa and CH for Switzerland.
Postal Code	Enter the postal code of the location in the site. For example, 100XX for New York.
State	Enter the state in which the site is located. If the country doesn't have states, enter the region in which the site is located. For example, NYC for New York City.
Devices	
Devices	Enter or select the CE devices to be used for the L3VPN service at the site. For example, ACX. This parameter allows the customer to request for devices from the service provider to be installed at the site.

Add L3VPN Site Network Access Parameters

Prerequisite:

["Add L3VPN Site" on page 777.](#)

After you add a site for the L3VPN service, you must configure the parameters for the network access (connection) between the site and the service provider network. A site network access defines how a site is connected to the service provider network.

To set site network access parameters:

1. Select and expand the site name in the **Customer Sites** table.
2. Click the add (+) icon above the Site Network Access table.

The Add Connection page appears.

3. Enter the site network access ID, device reference values, and connection type by referring to the following table:

Field	Description
Site Network Access ID*	Enter a unique ID for the connection between the site and the service provider network. For example, site-2.
Device Reference	Click the drop-down list to select a CE device for the connection. The drop-down list displays the devices that you specified on the Add Sites page.
Type*	Click the drop-down to select a connection type between the site and service provider network from the following options: • irb—Select irb to enable both Layer 2 bridging and Layer 3 routing on the same interface. • regular—Select regular for standard Layer 3 VPN connection type between the site and service provider network. By default, the regular connection type is selected. NOTE: • When you select irb as the connection type, the Device Reference, Access Diversity, Bearer, and Connection configuration parameters are hidden. • Configuring IRB interfaces for an L3VPN service is qualified only for the following scenarios: • L3VPN with EVPN having regular untagged interfaces with OSPF as PE-CE protocol. • L3VPN with EVPN having regular untagged interfaces with BGP as PE-CE protocol. • L3VPN with EVPN having interfaces in VLAN mode with OSPF as PE-CE protocol. • L3VPN with EVPN having interfaces in VLAN mode with BGP as PE-CE protocol.

4. Expand **Access Diversity** and add the access diversity parameters. See ["Add Access Diversity Parameters for L3VPN Site Network Access"](#) on page 786.

5. Expand **Bearer** and add the bearer parameters by referring to the following table:

Table 187: Bearer Parameters (L3VPN)

Field	Description
Bearer Reference	Enter the ID of an existing connection or bearer (access media and other L2 properties) between a service provider and customer, which the customer wants to reuse for the L3VPN service.
Requested Type	Select the access media type that the customer prefers to use. Ethernet is the only supported media type in this release.
Strict	Set to True if the requested access type is a strict requirement. If the requested access type is a strict requirement, the service provider cannot connect the site using any other media. The default setting is False .

6. Expand **Connection** and add the Ethernet interface parameters.

Click the **Ethernet interface type** drop-down and select the interface type from the following options:

- Untagged**—An untagged interface receives and sends Ethernet frames without VLAN tags. When selected, only unused devices and interfaces with no VLANs are assigned as placement options.
- Tagged**—A tagged interface receives and sends Ethernet frames with VLAN tags. When selected, unused devices and interfaces with the VLAN ID equal to the CVLAN ID configured here, are assigned as placement options.

Click the **Type** drop-down and select **dot1q** or **qinq** to define the type of tagged interface type.

- dot1q**—A Dot1q tagged interface transmits and receives Ethernet frames with VLAN tags. VLAN tags are added to the frames based on guidelines defined by the IEEE 802.1Q standard. In the Dot1q method, the service provider tags packets between CE and PE devices with one or more customer VLAN (CVLAN) IDs.

(Mandatory) Expand **Dot1q VLAN Tagged** and configure the CVLAN ID. Enter a unique identifier for the CVLAN. Enter a value in the range 1 through 4094.

- qinq**—A Q-in-Q interface transmits and receives frames with double VLAN tags that is an inner CVLAN and an outer service VLAN (SVLAN). In the Q-in-Q method, the service

provider adds an additional SVLAN tag outside the CVLAN tag, creating a double-tagged frame between the CE and PE device.

(Mandatory) Expand **Qinq** and configure CVLAN ID and SVLAN ID. Enter unique IDs for the inner CVLAN and the outer SVLAN. Enter a value in the range 1 through 4094.

7. Expand **IP Connection > Ipv4** and add the connection parameters by referring to the following table:

Table 188: IPv4 Connection Parameters (L3VPN)

Field	Description
Address Allocation Type	Define how to allocate IP addresses when you configure IPv4 subnet. The only supported type is static address, where IP addresses are assigned manually. NOTE: Ensure that the provider and customer addresses are in the same subnet. Also, if you select the static address allocation type, you must enter the provider and customer addresses with the prefix length.
Provider Address*	Enter the IPv4 address of the service provider for the customer network to connect, in the dotted decimal notation. For example, 10.10.3.4.
Customer Address*	Enter the IPv4 address of the customer network to connect with the service provider network, in the dotted decimal notation. For example, 192.168.1.2.
Prefix Length*	Enter the subnet prefix length expressed in number of bits. The same prefix length is applied to both customer and provider IPv4 addresses. For example, 20.

8. Expand **IP Connection > Ipv6** and add the connection parameters by referring to the following table:

Table 189: IPv6 Connection Parameters (L3VPN)

Field	Description
Address Allocation Type	Define how to allocate IP addresses when you configure IPv6 subnet. The only supported type is static address, where IP addresses are assigned manually. NOTE: Ensure that the provider and customer addresses are in the same subnet. Also, if you select the static address allocation type, you must enter the provider and customer addresses with the prefix length.
Provider Address*	Enter the IPv6 address of the service provider for the customer network to connect, in the hexadecimal format. For example, 2001:db8::1.
Customer Address*	Enter the IPv6 address of the customer network to connect with the service provider network, in the hexadecimal format. For example, 2001:db8::1::2.
Prefix Length*	Enter the subnet prefix length expressed in number of bits. The same prefix length is applied to both customer and provider IPv6 addresses. For example, 127.

You can configure only IPv4 or only IPv6 connection parameters or both.

- Expand **IRB** and add the virtual gateway details by referring to the following table:

Table 190: Virtual Gateway Parameters (L3VPN)

Field	Description
IP address	Enter the virtual gateway IP address assigned to the IRB interface. For example, 10.10.10.1. This IP address serves as the shared default gateway across PE devices and enables Layer 3 routing.

Table 190: Virtual Gateway Parameters (L3VPN) (Continued)

Field	Description
MAC address	Enter the virtual gateway MAC address assigned to the IRB interface. For example, 00:1a:2b:3c:4d:5e. This MAC address serves as the Layer 2 identifier for the virtual gateway and enables forwarding across PE devices.

**NOTE:**

- This section is available only when you select IRB as the connection type on the Add Connection page
- Configuring IRB interfaces for an L3VPN service is qualified only for specific scenarios.

10. Expand **Routing Protocols** and add the routing protocols. See ["Add Routing Protocol for L3VPN Site Network Access"](#) on page 788.
11. Expand **Service** and add the service parameters by referring to the following table:

Table 191: Service Parameters (L3VPN)

Field	Description
Service Input Bandwidth	Enter the download bandwidth (in bps) for the site from the provider network. Enter a value from 0 to 9,223,372,036,854,766,000 bps.
Service MTU	Enter the maximum packet size (in bytes) allowed through the connection for the L3VPN service. Enter a value from 0 to 65,535 bytes.
Service Output Bandwidth	Enter the upload bandwidth (in bps) from the site to the service provider network. Enter a value from 0 to 9,223,372,036,854,766,000 bps.

12. Expand **VPN Attachment** > **Attachment Flavor** and click the **Site Role** drop-down to select the role of the site in the VPN topology:

- Any-to-any role—In this role, all VPN sites can communicate with each other.
- Hub role—The site has a hub role in a hub-and-spoke VPN service topology.
- Spoke role—The site has a spoke role in a hub-and-spoke VPN service topology.



NOTE: The any-to-any role is supported if you select the any-to-any VPN service topology for the L3VPN service. The hub-and-spoke roles are supported if you select the hub-spoke service topology for the L3VPN service.

13. Configure placement parameters for the service instance:

- Click the **Update Placements** button on top of the page and click **Yes** when prompted to confirm. The service instance is saved to the Routing Director database. The saved instance is listed in the Service Instances table on the Service Instances page.
Service orchestration automatically assigns possible placement options such as PE devices and interfaces, CE devices, and VLANs for the service instance. After placement options are assigned, you see a message that placement was successfully updated.
- Proceed to the Site Settings section, select the site network access for which you want to configure placement, and click the edit icon.
The Edit Connection page appears and the **Placement** section is displayed on the page.
- Expand **Placement** or **Placement IRB** and configure the placement parameters.
 - If you selected **regular** as the connection type on the Add Connection page, configure the placement parameters by referring to the following table:

Table 192: Placement Parameters

Field	Description
CE Node	Click the CE Node drop-down and select the CE device that you want to connect to the PE device for the VPN service.
PE Node	Click the PE Node drop-down and select the PE device that you want to allocate for the VPN service.

Table 192: Placement Parameters *(Continued)*

Field	Description
PE Interface	Click the PE Interface drop-down and select the PE interface name that you want to assign for the VPN service.

- If you selected **irb** as the connection type on the Add Connection page, configure the placement parameters by referring to the following table:

Table 193: Placement IRB Parameters

Field	Description
Instance	Select the service instance that created the IRB interface for L3VPN service.
NE Id	Select the network element ID that identifies the device to which the IRB interface belongs.
Name	Displays the name of the interface assigned to the IRB instance.



NOTE: Configuring IRB interfaces for an L3VPN service is qualified only for specific scenarios.

14. Do any of the following:

- Click **Cancel** to exit the Add Connection page without saving the changes you made.
- Click **OK** to save the connection details you added.

You are returned to the Add L3VPN Service page and the connection you added is listed in the Site Network Access table.

Add Access Diversity Parameters for L3VPN Site Network Access

Prerequisite:

Complete steps 1 to 3 of the ["Add L3VPN Site Network Access Parameters" on page 779](#) procedure.

A Network Admin can group the connections from a site to the service provider network, assign group ID to each connection group, and apply certain constraints to all the connections included in a group. You can assign one or more group IDs for a connection, specify constraint types, and select the target group ID that the constraint must be applied to, in the Access Diversity section.

To assign group IDs to a connection and apply constraints to target group IDs:

1. Expand **Access Diversity** on the Add Connection page.
2. Add IDs of groups that the connection is to be a part of:
 - a. Expand **Groups** and click the add (+) icon.
The Group page appears.
 - b. Enter one or more IDs of the groups to which the connection must be assigned and click **OK**.
The group ID entries are listed in the **Group ID** table.
 - c. (Optional) Use the edit or delete options above the Group table to edit or delete the entries.
3. Expand **Constraints** and click the add (+) icon above the Constraints table.
The Constraint page appears.
4. Add the following parameters on the Constraints page:
 - a. Click the **Constraint Type** drop-down and select the constraint type from the following supported options:
 - **pe-diverse**—The site network access must not be connected to the same PE router as the target connections from the site.
 - **same-pe**—The site network access must be connected to the same PE router as the target connections from the site.
 - **pop-diverse**—The site network access must not use the same point-of-presence (POP) as the target connections from the site.
 - b. Expand **Target** and click the add (+) icon above the Group ID table.
The Group page appears.
 - c. Click the Group ID drop-down to select the group to which you want to apply the constraint and click **OK**.
The ID you enter is displayed in the Group ID table.

You can add multiple group IDs as target groups to apply the access constraint.
 - d. (Optional) Use the edit or delete options above the Constraints table to edit or delete the entries.
 - e. Click **OK** on the Constraint page.
The constraint types and target group IDs are listed in the **Constraint** table.

5. Proceed to step 5 of the "Add L3VPN Site Network Access Parameters" on page 779 procedure to add the remaining site network access parameters.

Add Routing Protocol for L3VPN Site Network Access

Prerequisite:

Complete steps 1 to 7 of the "Add L3VPN Site Network Access Parameters" on page 779 procedure.

The Routing Protocols section defines the routing protocol to be used between the PE and CE routers.

In a customer-managed site, the routing protocol that you define here is used between the PE router and CE router that the customer manages. In a service provider-managed site, the routing protocol that you define here is used between the provider-managed CE router and the customer-managed LAN or router. Routing Director supports the static route, direct, OSPF, and BGP protocols for L3VPN service provisioning.

To set routing protocols:

1. Expand **Routing Protocols** on the Add Connection page.
2. Click the Add button present above the Routing Protocols table and set the routing protocol from the supported options:



NOTE: The hub-and-spoke topology supports only the static route protocol.

- Static route—To set the static route protocol:
 - a. Expand **Routing Protocols** and click **Add > Static Route** on the Add Connection page.
The Add Static Route page appears.
 - b. Expand **Cascaded LAN Prefixes** and click the add (+) icon next to Ipv4 LAN Prefixes.
The IPv4 LAN Prefixes page appears.
 - c. Enter the values as described in the following table:

Table 194: IPv4 Static Route Parameters (L3VPN)

Field	Description
LAN	Enter the IPv4 address with prefix from the customer LAN to be routed through the L3VPN connection. For example, 192.168.0.1/20.

Table 194: IPv4 Static Route Parameters (L3VPN) (Continued)

Field	Description
Next Hop	Enter the IPv4 address of the next hop router to be used for forwarding packets in the static route. For example, 10.1.0.1.



NOTE: In the L3VPN hub-and-spoke topology, by default, the spokes cannot communicate with each other and they can only reach the hub CE device. To ensure that spokes communicate with each other through the hub CE device, add a 0.0.0.0/0 LAN interface with the customer device interface IP address as the next hop.

- d. Click **OK**.

The static route details are listed in the IPv4 LAN Prefixes table.

- e. Expand **Cascaded LAN Prefixes** and click the add (+) icon next to Ipv6 LAN Prefixes.

The IPv6 LAN Prefixes page appears.

- f. Enter the values as described in the following table:

Table 195: IPv6 Static Route Parameters (L3VPN)

Field	Description
LAN	Enter the IPv6 address with prefix from the customer LAN to be routed through the L3VPN connection. For example, 2001:db8::1:2/127.
Next Hop	Enter the IPv6 address of the next hop router to be used for forwarding packets in the static route. For example, 2001:db8::2:1.

- g. Click **OK**.

The static route details are listed in the IPv6 LAN Prefixes table.

- h. (Optional) To edit or delete an entry, select the entry and click the edit or delete icons present above the IPv4 or IPv6 LAN Prefixes tables.

- OSPF—To set the OSPF routing protocol:
 - a. Expand **Routing Protocols** and click **Add > OSPF** on the Add Connection page.

The Add OSPF page appears.

- b. Enter the values as described in the following table:

Table 196: OSPF Parameters (L3VPN)

Field	Description
Address Family	Select the IP addressing protocol to be configured for the OSPF sessions. You can select either IPv4, IPv6, or both.
Area Address*	Enter the OSPF area address to be used for the interface connecting to the service provider network. For example, 0.0.0.1.
Metric	Enter the cost of using the OSPF link to the service provider network. Enter a number in the range 1 through 65,535.

- c. Click **OK**.

The OSPF details are listed in the Routing Protocols table.

- BGP—To set the BGP routing protocol:
 - a. Expand **Routing Protocols** and click **Add > BGP** on the Add Connection page.

The Add BGP page appears.

 - b. In the **Address Family** field, select the IP addressing protocol to be configured for the BGP sessions. You can select either IPv4, IPv6, or both.
 - c. (Mandatory) In the **Autonomous System** field, enter the autonomous system number of the network in which you want to configure the L3VPN service.
 - d. Click **OK**.

The BGP details are listed in the Routing Protocols table.

- Direct—Use the direct routing protocol when a customer LAN at the site is directly connected to the service provider network.

To use the direct routing, click **Add > Direct**. The entry is listed in the Routing Protocols table.

3. (Optional) To edit or delete an entry, select the entry and click the edit or delete icons present above the Routing Protocols table.
4. Proceed to step 11 of the ["Add L3VPN Site Network Access Parameters" on page 779](#) procedure to add the remaining site network access parameters.

SEE ALSO

[Add an L3VPN Service Instance | 773](#)

[Modify a Service Instance | 822](#)

Add an EVPN or EVPN-VPWS Service Instance

A superuser or network administrator can use Routing Director to provision an L2VPN service in their network. Routing Director supports provisioning of Ethernet VPN (EVPN) and EVPN Virtual Private Wire Service (VPWS) services.

An EVPN service provides multipoint Layer 2 connectivity and enables you to connect dispersed customer sites or devices, while an EVPN-VPWS service provides point-to-point connection between two customer sites or devices over an MPLS network.

When you create and save an EVPN or EVPN-VPWS service instance, Routing Director generates a create service order. After you provision the service instance, Routing Director activates the automated workflow for fulfilling the service order and provisions the service in the network.

You can create EVPN and EVPN-VPWS service instances by uploading preconfigured JSON files or by entering the details in the fields on the Add E-LAN EVPN CSM and Add E-Line EVPN VPWS CSM service pages respectively.

To create an EVPN or an EVPN-VPWS service instance:

1. Click **Orchestration > Instances**.
The Service Instances page appears.
2. To add an EVPN service instance, click **Add > E-LAN EVPN CSM**. The Add E-LAN EVPN CSM service page appears.
To add an EVPN-VPWS service instance, click **Add > E-Line EVPN VPWS CSM**. The Add E-Line EVPN VPWS CSM service page appears.
3. In the General section of the Add E-LAN EVPN CSM page, enter the values by referring to [Table 197 on page 793](#).

In the General section of the Add E-Line EVPN VPWS CSM page, enter the values by referring to [Table 198 on page 795](#).

4. Click **Next** to proceed to the **Site Settings** section.
5. Enter site-specific details in the Site Settings section. See ["Add EVPN or EVPN-VPWS Service Site Details" on page 796](#).

6. Click **Next** to proceed to the **Summary** section.

You see a graphical map of your service topology and a summary of the instance details in the **Summary** section. You can zoom in and zoom out using the zoom options on the map.

7. (Optional) Click **Export** to export and save the service details in JSON format on your local system.
8. (Optional) Click **Reset Placements** to remove existing placement configurations for the service instance.

Provisioning of a service instance might fail when existing placement configurations for the service instance are no longer available or applicable. When you reset placement parameters, service orchestration automatically assigns placement options for the instance before provisioning the service.

9. Click **Next** to proceed to the Schedule section.
10. (Optional) Switch the Schedule toggle on if you want to schedule when to provision the service. The Schedule option is disabled by default. When you schedule provisioning, the **Save & Provision** button on top of the page is disabled and grayed out.

The Schedule field appears. Perform the following steps to schedule the date and time to provision the service:

- a. Click the calendar icon in the Schedule field.
- b. Select the date and time for provisioning the service, and click **OK**.
The scheduled date and time are displayed in the MM/DD/YYYY HH:MM AM/PM format.



NOTE:

- You must set a present or future date and time to schedule provisioning of a service. If you select a past time and save the service order, the service is provisioned immediately.
- To delete a provisioning schedule, select the *Service-Instance-Name* on the Service Instances page and click the edit icon. On the Modify *Service-Instance-Name* page, proceed to the Schedule section. In the Schedule section, switch the Schedule toggle off and click **Save**. The provisioning schedule is deleted and the service instance is provisioned immediately.

- Scheduling of service provisioning is a Beta feature in this release.

11. To save and provision the service instance, click one of the following buttons:

- **Save & Provision**—Click this button and confirm when prompted, to save the service instance to the Routing Director database and provision the instance immediately. This option is disabled if you have scheduled the provisioning of a service instance.
- **Save**—Click this button to save the service instance to the Routing Director database.

The Service Instances page appears and the new service instance is listed in the Service Instances table. To provision the service instance, select the service instance and click **Provision**. Alternatively, you can schedule provisioning of the service instance as defined in step 10.



NOTE:

- The **Provision** button is disabled if you schedule the provisioning of a service instance. The service is provisioned automatically at the scheduled date and time.



NOTE: Fields marked with an asterisk (*) are mandatory.

Table 197: Fields in the Add E-LAN EVPN CSM Service Page General Section

Field	Description
Upload JSON File	Click Browse to upload a preconfigured JSON file. You see a message that the file is successfully imported. The values specified in the file are automatically populated in the corresponding UI fields. NOTE: If you are uploading a preconfigured JSON file to create an EVPN service instance, you must clear the placement section in the file before you provision the service order.

Table 197: Fields in the Add E-LAN EVPN CSM Service Page General Section (*Continued*)

Field	Description
Customer*	Enter the name of the customer for whom you are provisioning the service: - If you already added the customer name by using the Customer Inventory page (Orchestration > Service > Customers), click the Customer field to see a drop-down with the customer names that you added. Select the customer name from the drop-down list. - Alternatively, click the Add Customer link to create a new customer. See "Add a Customer" on page 666. The customer name must be unique within an organization. For example, network-operator.
Instance name*	Enter a name for the service instance. For example, evpn1. The instance name can be a set of alphanumeric characters and the special character hyphen (-). The maximum number of characters allowed is 64.
VPN Id*	Enter the ID you want to assign to the VPN. For example, vpn10. The VPN ID must not exceed 64 characters.
VPN Service Topology	The topology for the VPN service. Only the any-to-any topology is supported in this release. In this service topology, all VPN sites can exchange network traffic with each other without any restrictions.
VPN Service Type	The service type for EVPN service provisioned by the service provider. Routing Director supports only the EVPN-MPLS service type.

Table 198: Fields in the Add E-LINE EVPN VPWS CSM Service Page General Section

Field	Description
Upload JSON File	Click Browse to upload a preconfigured JSON file. You see a message that the file is successfully imported. The values specified in the file are automatically populated in the corresponding UI fields. NOTE: If you are uploading a preconfigured JSON file to create an EVPN-VPWS service instance, you must clear the placement section in the file before you provision the service order.
Customer*	Enter the name of the customer for whom you are provisioning the service: • If you already added the customer name by using the Customer Inventory page (Orchestration > Service > Customers), click the Customer field to see a drop-down with the customer names that you added. Select the customer name from the drop-down list. • Alternatively, click the Add Customer link to create a new customer. See "Add a Customer" on page 666. The customer name must be unique within an organization. For example, network-operator.
Instance name*	Enter a name for the service instance. For example, evpn-vpws1. The instance name can be a set of alphanumeric characters and the special character hyphen (-). The maximum number of characters allowed is 64.

RELATED DOCUMENTATION

[Service Provisioning Workflow | 648](#)

[Add EVPN or EVPN-VPWS Service Site Details | 796](#)

[Modify a Service Instance | 822](#)

Add EVPN or EVPN-VPWS Service Site Details

IN THIS SECTION

- [Add EVPN or EVPN-VPWS Site | 796](#)
- [Add EVPN or EVPN-VPWS Site Network Access Parameters | 799](#)
- [Add Access Diversity Parameters for EVPN and EVPN VPWS Site Network Access | 804](#)
- [Add Connection Parameters for EVPN or EVPN-VPWS Site Network Access | 806](#)

An EVPN service is provisioned on a collection of sites to exchange network traffic over a shared infrastructure with a common layer 2 technology such as Ethernet. An EVPN-VPWS service is provisioned on a point-to-point connection between two sites over an MPLS network. Use the Customer Site Settings section of the Add E-LAN EVPN CSM or the Add E-Line EVPN VPWS CSM service page to enter details of sites that you want to connect to the EVPN or EVPN-VPWS.

You configure the following in the Customer Site Settings section:

- General details (site ID, site VPN flavor, location, devices, and MAC loop prevention parameters) about the site. See ["Add EVPN or EVPN-VPWS Site" on page 796](#).
- Parameters for site network access (connection with the service provider network). See ["Add EVPN or EVPN-VPWS Site Network Access Parameters" on page 799](#).

Add EVPN or EVPN-VPWS Site

Prerequisite

Before adding the EVPN or EVPN-VPWS service site details, you must create a new EVPN or EVPN-VPWS service instance and enter general details about the service. Follow steps 1 to 4 of the ["Add an EVPN or EVPN-VPWS Service Instance" on page 791](#) procedure to create a new instance and add general details.

A site refers to the geographical area where the VPN service spans. To add an EVPN or EVPN-VPWS service site in the Customer Site Settings section:

1. Click the add (+) icon on the top-right corner of the Customer Sites section.
The Add Site page appears.
2. Enter site details by referring to [Table 199 on page 797](#).
3. Do any of the following:
 - Click **Cancel** to exit the Add Site page without saving the changes you made.

- Click **OK** to save the site details you added.

The site you added is listed in the **Sites** table.

You can view the configured site properties by expanding *site-ID* > **Properties**.

Configure the network access parameters. See ["Add EVPN or EVPN-VPWS Site Network Access Parameters" on page 799](#).



NOTE: Fields marked with an asterisk (*) are mandatory.

Table 199: Fields on the Add Site Page

Field	Description
Site ID*	Enter a unique site ID to identify the site in the network. For example, site1. The site ID can contain alphanumeric characters and special characters hyphens (-) and underscores (_). The maximum number of characters allowed is 64.
Site VPN Flavor	Select how you want to connect the site to the VPN. For example, a site can connect to a single VPN or multiple VPNs. Only single site VPN flavor (site-vpn-flavor-single) is supported in this release. In the single site VPN flavor, the site is connected to only one VPN. NOTE: This field is available only when you create an E-LAN EVPN CSM service instance.
Locations	
Location ID*	Enter a unique ID for a location in the site. For example, MAN for Manhattan and BRO for Brooklyn in the New York site.
Address	Enter the address (number and street) of the location in a site. For example, 123, Avenue Street.

Table 199: Fields on the Add Site Page *(Continued)*

Field	Description
City	Enter the city where the site is located. For example, MAN for Manhattan.
Country Code	Enter the ISO alpha-2 code of the country where the site is located, for example, ZA for South Africa and CH for Switzerland.
Postal Code	Enter the postal code of the location in the site. For example, 100XX for New York.
State	Enter the state in which the site is located. If the country doesn't have states, enter the region in which the site is located. For example, NYC for New York City.
Devices	
Devices	Enter or select the CE devices to be used for the EVPN or EVPN-VPWS service at the site. For example, ACX-CE1. This parameter allows the customer to request for devices from the service provider to be installed at the site.
Expand MAC Loop Prevention This section is available only when you create an E-LAN EVPN CSM service instance.	
Frequency	Enter the number of times that MAC address duplication is to be detected in a default time interval of 180 seconds or 3 minutes. When a MAC address is detected as duplicate, it is added to a list of duplicate MAC addresses. The default frequency is 5.

Table 199: Fields on the Add Site Page (*Continued*)

Field	Description
Retry Timer	Set the retry interval (in seconds) for a duplicate MAC address. Enter a value in the range 60 to 21600. After the retry timer expires, the duplicate MAC address is removed from the MAC-VRF table (virtual routing and forwarding table for MAC addresses) on PE devices.

Add EVPN or EVPN-VPWS Site Network Access Parameters

Prerequisite:

["Add EVPN or EVPN-VPWS Site" on page 796.](#)

After you add a site for the EVPN or EVPN-VPWS service, you must configure the parameters for site network access (connection) between the site and the service provider network. Site network access refers to a logical Ethernet connection from a site to the L2VPN service in the provider network.

To set site network access parameters:

1. Select and expand the *site-ID* in the **Customer Sites** table.
2. Expand **Site Network Access** and click the add (+) icon above the Site Network Access table.
The Add Connection page appears.
3. (Mandatory) In the **Network Access ID** field, enter a unique ID for the connection between the site and the service provider network. For example, SVL-access-1.
4. In the **Device Reference** field, click the drop-down list to select a CE device for the connection. The drop-down list displays the devices that you specified on the Add Sites page.
5. Expand **Access Diversity** and add the access diversity parameters. See ["Add Access Diversity Parameters for EVPN and EVPN VPWS Site Network Access" on page 804.](#)
6. Expand **Bearer** and add the bearer parameters by referring to the following table:

Table 200: Bearer Parameters

Field	Description
Strict	Set to True if the requested access type is a strict requirement. If the requested access type is a strict requirement, the service provider cannot connect the site using any other media. The default setting is False .
Type*	Specify the access media type that the customer prefers to use. Ethernet is the only supported media type in this release.



NOTE: This section is available only when you create an E-LAN EVPN CSM service instance.

7. Expand **Connection** and add the connection parameters. See ["Add Connection Parameters for EVPN or EVPN-VPWS Site Network Access" on page 806](#).
8. Expand **Availability** and enter the values by referring to the following table:

Table 201: Availability Parameters

Field	Description
Access Priority	Enter a value for site network access priority. Enter a value in the range 0 through 4,294,967,295. The default value is 100. The access priority parameter defines the preference for the site network access (connection) for load balancing and configuring primary or backup site connections. The higher the access priority value, the higher is the preference for the connection.
Redundancy Mode	Select the redundancy mode from the following options: • all-active—In this mode, all nodes can forward traffic to and from the Ethernet segment. • single-active—In this mode, only one node forwards traffic to and from the Ethernet segment.

9. Add service parameters for the site network access:
 - a. Expand **Service** > **Service Bandwidth** and click the add (+) icon above the Bandwidth table.
The Bandwidth page appears.

- b. Enter the values by referring to the following table:

Table 202: Service Parameters

Field	Description
Direction	Select the direction of service bandwidth: • input-bw—Bandwidth download direction from the service provider network to the site. • output-bw—Bandwidth upload direction from the site to the service provider network.
CBS*	Enter the Committed Burst Size (CBS) in bytes per second. When network traffic does not utilize the configured Committed Information Rate (CIR), the unused bandwidth is accumulated up to a maximum limit defined by the CBS.
CIR*	Enter the CIR or the maximum number of bits per second (bps) that an interface can send or receive. The service provider provides the CIR as the guaranteed bandwidth for network traffic that an interface can send or receive in one second, under normal line conditions.

- c. Click **OK**.
The bandwidth values are listed in the Bandwidth table.
 - d. (Optional) Click the edit or delete icons above the Bandwidth table to edit or delete the entries.
10. For E-LAN EVPN CSM service instances, expand **Layer3 Gateway** and switch the **Enabled** toggle on to set it to **True** to enable an integrated routing and bridging (IRB) interface on a PE device. IRB supports Layer 2 bridging and Layer 3 routing on the same interface.
The toggle button is set to **False** and disabled by default.



NOTE:

- This section is available only when you create an E-LAN EVPN CSM service instance.
- Configuring IRB interfaces for an E-LAN EVPN CSM service is qualified only for the following scenarios:
 - EVPN with Ethernet interface port mode.
 - EVPN with Ethernet interface VLAN mode.

11. Expand **MAC address limit** and enter the values by referring to the following table:

Table 203: MAC Address Limit Parameters

Field	Description
Action	Specify the action to be taken when the upper limit for customer MAC addresses is exceeded. The service provider specifies a maximum number of MAC addresses that can be learned from a customer for a single service instance. Drop is the only action supported in this release. This action drops the packet when the maximum limit for MAC addresses is exceeded.
Limit Number	Specify the maximum number of MAC addresses that can be obtained from a customer for a single service instance. Enter a value in the range 0 and 4,294,967,295.



NOTE: This section is available only when you create an E-LAN EVPN CSM service instance.

12. Configure placement parameters for the service instance:

- Expand **Placement** and configure the placement parameters by referring to the following table:

Table 204: Placement Parameters

Field	Description
CE Node	Click the CE Node drop-down and select the CE device that you want to connect to the PE device for the VPN service.

Table 204: Placement Parameters *(Continued)*

Field	Description
PE Node	Click the PE Node drop-down and select the PE device that you want to allocate for the VPN service.
PE Interface	Click the PE Interface drop-down and select the PE interface name that you want to assign for the VPN service.
Speed	The speed that you entered while configuring the topology resource instance is displayed here in Mbps.

- b. Alternatively, you can update placements by exiting this wizard by clicking **OK**. Click the **Update Placements** button on top of the page and click **Yes** when prompted to confirm.

The service instance is saved to the Routing Director database. The saved instance is listed in the Service Instances table on the Service Instances page.

Service orchestration automatically assigns possible placement options such as PE devices and interfaces, CE devices, and VLANs for the service instance. After placement options are assigned, you see a message that placement was successfully updated.

- 13.** Expand **Placement IRB** and configure the placement parameters.

The **Placement IRB** section appears when you click the **Update Placements** button on the top of the page. Click **Yes** to confirm.

Configure the placements IRB parameters by referring to the following table:

Table 205: Placement IRB Parameters

Field	Description
POP	Click the POP drop-down and select the point-of-presence (POP) site that hosts the IRB interface. A POP refers to a site or physical location that hosts the IRB interface.

Table 205: Placement IRB Parameters *(Continued)*

Field	Description
PE	Click the PE drop-down and select the PE device that you want to assign for the IRB interface.
IRB	Displays the IRB interface ID to be configured on the PE device. The range of IRB interface IDs that you entered while configuring the topology resource instance is displayed next to the IRB text box.



NOTE: Configuring IRB interfaces for an E-LAN EVPN CSM service is qualified only for specific scenarios.

14. Do any of the following:

- Click **Cancel** to exit the Add Connection page without saving the changes you made.
- Click **Reset Placements** to delete existing placement configurations for the site and delete the **Placements** section from the page.
- Click **OK** to save the connection details you added.

You are returned to the Add E-LAN EVPN CSM service page after you save the EVPN connection details. You are returned to the Add E-Line EVPN VPWS CSM service page after you save the EVPN-VPWS connection details. The connection you added is listed in the Site Network Access table.

- 15.** (Optional) Click the edit or delete icons above the Site Network Access table to edit or delete the entries.
- 16.** Go to step 6 of the ["Add an EVPN or EVPN-VPWS Service Instance" on page 791](#) procedure to proceed with the EVPN or EVPN-VPWS service instance creation or modification.

Add Access Diversity Parameters for EVPN and EVPN VPWS Site Network Access



NOTE: This section is available only when you create an E-LAN EVPN CSM service instance.

Prerequisite:

Complete steps 1 to 3 of the ["Add EVPN or EVPN-VPWS Site Network Access Parameters" on page 799](#) procedure.

A network administrator can group the connections from a site to the service provider network, assign group ID to each connection group, and apply certain constraints to all the connections included in a group. A constraint is expressed in relation to another set of connections (target set). For example, a connection X in a group with group ID G1 must be "PE-diverse" (constraint type) from connections in a target group with group ID G2. The site network access that you are configuring might use the same or different group ID than the target network accesses. You can assign one or more group IDs for a connection, specify constraint types, and select the target group ID that the constraint must be applied to, in the Access Diversity section.

To assign group IDs to a connection and apply constraints to target group IDs:

1. Expand **Access Diversity** on the Add Connection page.
2. Add IDs of groups that the connection is to be a part of:
 - a. Expand **Groups** and click the add (+) icon.
The Group page appears.
 - b. Enter one or more IDs of the groups to which the connection must be assigned and click **OK**.
The group ID entries are listed in the **Group ID** table.
 - c. (Optional) Use the edit or delete options above the Group table to edit or delete the entries.
3. Expand **Constraints** and click the add (+) icon above the Constraint table.
The Constraint page appears.
4. Add the following parameters on the Constraints page:
 - a. Click the **Constraint Type** drop-down and select constraint type from the following supported options:
 - **pe-diverse**—The site network access (connection) must not be connected to the same PE router as the connections in the target group.
 - **pop-diverse**—The site network access must not use the same point-of-presence (POP) as the connections in the target group.
 - **same-pe**—The site network access must be connected to the same PE router as the connections in the target group.



NOTE: The **same-pe** constraint type is available only for E-LAN EVPN CSM service instances.

- b. Expand **Groups** and click the add (+) icon above the Group ID table.
The Group page appears.

- c. Click the **Group ID** drop-down to select the group to which you want to apply the constraint and click **OK**.

The ID you enter is displayed in the Group ID table.

You can add multiple group IDs as target groups to apply the access constraint.

- d. Click **OK** on the Constraint page.

The constraint types and target group IDs are listed in the **Constraint** table.

- e. (Optional) Click the edit or delete icons above the Constraints table to edit or delete the entries.

5. Go to step 6 of the ["Add EVPN or EVPN-VPWS Site Network Access Parameters" on page 799](#) procedure to add the remaining site network access parameters.

Add Connection Parameters for EVPN or EVPN-VPWS Site Network Access

Prerequisite:

Complete steps 1 to 6 of the ["Add EVPN or EVPN-VPWS Site Network Access Parameters" on page 799](#) procedure to add connection parameters for the EVPN site network access.

Complete steps 1 to 4 of the ["Add EVPN or EVPN-VPWS Site Network Access Parameters" on page 799](#) procedure to add connection parameters for the EVPN-VPWS site network access.

To add connection parameters:

1. Expand **Connection** on the Add Connection page.
2. Click the Ethernet interface type drop-down and select the interface type from the following options:
 - Tagged—A tagged interface receives and sends Ethernet frames with VLAN tags.
 - Untagged—An untagged interface receives and sends Ethernet frames without VLAN tags.
3. For E-LAN EVPN CSM service instances, expand **Lag Interfaces** and enter the Link Aggregation Group (LAG) interface configurations:
 - a. Click the add (+) icon above the Lag Interface table.
The Lag Interface page appears.
 - b. In the **Index** field, enter a unique ID to identify the LAG in the network.
The customer can aggregate multiple links from physical interfaces on a CE device into a single, logical LAG interface to connect to the service provider network.
 - c. Expand **LACP** and enter the values by referring to the following table:

Table 206: LACP Parameters

Field	Description
Enabled	Switch the Enabled toggle on to set it to True if you want to enable the Link Aggregation Control Protocol (LACP) between CE and PE devices to dynamically manage adding or deleting LAG member links. The default setting is False .
LLDP	Switch the LLDP toggle on to set it to True if you want to enable Link Layer Discovery Protocol (LLDP). The default setting is False .
Minimum number of links	Enter the minimum number of links that must be active to use the LAG link to connect to the network. Enter a value in the range 1 through 64.



NOTE: The Lag Interfaces section is available only when you are create an E-LAN EVPN CSM service instance.

- d. Expand **Member Links** and click the add (+) icon above the Member Link table.
The Member Link page appears.
- e. Enter the member link values by referring to the following table:

Table 207: LACP Member Link Parameters

Field	Description
Index	Enter a unique ID to identify the member link.
Speed	Enter the speed (in Mbps) of the physical interface or member link that is to be part of the LAG. Enter a value in the range 0 through 4,294,967,295, for example, 10 Mbps.

- f. Expand **OAM 802.3ah Link**.
- g. Switch the **Enabled** toggle on to set it to **True**, if you want to enable IEEE 802.3ah Operation, Administration, and Management (OAM) link.
The default setting is **False**.
- h. Click **OK** on the Member Link page.
The member link configurations are listed in the Member Link table.

Alternatively, click **Cancel** to exit the Member Link page without saving the changes you made.
- i. (Optional) Click the edit or delete icons above the Member Link table to edit or delete the entries.
- j. Click **OK** on the Lag Interface page.
Alternatively, click **Cancel** to exit the Lag Interface page without saving the changes you made.
- k. (Optional) Click the edit or delete icons above the Lag Interface table to edit or delete the entries.

4. (For tagged interface type only) Expand **Tagged Interface** and enter the interface parameters:

- a. For EVPN-VPWS services, click the **Type** drop-down and select **dot1q** to define the type of tagged interface.

Only the Dot1q tagged interface type is supported in this release. A Dot1q tagged interface transmits and receives Ethernet frames with VLAN tags. VLAN tags are added to the frames based on guidelines defined by the IEEE 802.1Q standard. In the Dot1q method, the service provider tags packets between CE and PE devices with one or more customer VLAN (CVLAN) IDs.

Select **dot1q** from the **Type** drop-down and configure the values by referring to the following table:

Table 208: Dot1q VLAN Tagged Parameters on the Edit Connection Page

Field	Description
CVLAN ID*	Enter a unique identifier for the CVLAN. Enter a value in the range 1 through 4094.
Tag type	Specify the type of VLAN tag. Click the Tag type drop-down and select c-vlan . Only the c-vlan tag type is supported this release.

- b. For EVPN services, click the **Type** drop-down and select **dot1q** or **qinq** to define the type of tagged interface.

- Select **dot1q** from the **Type** drop-down and configure the values by referring to the following table:

Table 209: Dot1q VLAN Tagged Parameters on the Edit Connection Page

Field	Description
CVLAN ID*	Enter a unique identifier for the CVLAN. Enter a value in the range 1 through 4094.
Tag type	Specify the type of VLAN tag. Click the Tag type drop-down and select c-vlan . Only the c-vlan tag type is supported this release.

- A Q-in-Q tagged interface transmits and receives frames with double VLAN tags—an inner CVLAN and an outer Service VLAN (SVLAN). In the Q-in-Q method, the service provider adds an additional SVLAN tag outside the CVLAN tag, creating a double-tagged frame between the CE and PE device.

Select **Qinq** from the **Type** drop-down and configure the values by referring to the following table:

Table 210: Q-in-Q VLAN Tagged Parameters on the Edit Connection Page

Field	Description
Cvlan Id*	Enter a unique identifier for the inner VLAN. Enter a value in the range 1 through 4094.
Svlan Id*	Enter a unique identifier for the outer VLAN. Enter a value in the range 1 through 4094.

5. (For untagged interface type only) For EVPN services, expand **Physical Interface Properties** on the Add Connection page of the EVPN service and configure the values by referring to the following table:

For EVPN-VPWS services, expand **Untagged Interface** on the Add Connection page of the EVPN-VPWS service and configure the values by referring to the following table:

Table 211: Physical Interface Properties and Untagged Interface Properties

Field	Description
LLDP	Switch the toggle on to set LLDP to True if you want to enable LLDP. The default setting is False .
Speed	Enter the speed (in Mbps) of the untagged physical interface. The default value is 10 Mbps.
Expand OAM 802.3ah Link	
Enabled	Switch the Enabled toggle on to set it to True , if you want to enable IEEE 802.3ah OAM link. The default setting is False .

6. Go to step 8 of the ["Add EVPN or EVPN-VPWS Site Network Access Parameters" on page 799](#) procedure to add the remaining site network access parameters for the E-LAN EVPN CSM service instance.

Go to step 9 of the ["Add EVPN or EVPN-VPWS Site Network Access Parameters" on page 799](#) procedure to add the remaining site network access parameters for the EVPN-VPWS service instance.

SEE ALSO

[Add an EVPN or EVPN-VPWS Service Instance | 791](#)

[Modify a Service Instance | 822](#)

Add an L2 Circuit Service Instance

As a superuser or network administrator, you can use Routing Director to provision Layer 2 (L2) circuit services in your network. An L2 circuit is a virtual point-to-point L2 connection between two CE devices through a packet-switched network such as Multiprotocol Label Switching (MPLS), in the service provider's network. The CE devices running the same L2 protocol (such as Ethernet) connect to PE devices in the service provider network through a physical or logical attachment circuit.

**NOTE:**

- LDP signaling without BGP autodiscovery is the only variant of L2 Circuit that is supported in this release.
- Redundant pseudowires for L2 Circuit are not supported in this release.

When you create and save an L2 circuit service instance, Routing Director generates a create service order. After you provision the service instance, Routing Director activates the automated workflow for fulfilling the service order and provisions the service in the network.

To create an L2 circuit service instance, either upload a preconfigured JSON file or enter the details in the GUI fields on the Add E-Line L2Circuit NSM page.

To create an L2 circuit service instance:

1. Click **Orchestration > Instances**.
The Service Instances page appears.
2. Click **+ > E-Line L2Circuit NSM**.
The Add E-Line L2Circuit NSM page appears.
3. In the General section, enter the values by referring to [Table 212 on page 813](#).
4. Click **Next** to proceed to the VPN Nodes section.
5. Enter VPN node details in the VPN Nodes section. See ["Add L2 Circuit VPN Nodes" on page 814](#).
6. Click **Next** to proceed to the Summary section.
You see a graphical map of your service topology and a summary of the instance details. Zoom in and out using the zoom icons on the map.
7. (Optional) Click **Export** to export and save the service details in JSON format on your local system.
8. (Optional) Click **Reset Placements** to remove existing placement configurations for the service instance.

Provisioning of a service instance might fail when existing placement configurations for the service instance are no longer available or applicable. When you reset placement parameters, service orchestration automatically assigns placement options for the instance before provisioning the service.
9. Click **Next** to proceed to the Schedule section.
10. (Optional) Switch the Schedule toggle on if you want to schedule when to provision the service.
The Schedule option is disabled by default. When you schedule provisioning, the **Save & Provision** button on top of the page is disabled and grayed out.

The Schedule field appears. Perform the following steps to schedule the date and time to provision the service:
 - a. Click the calendar icon in the Schedule field.

- b. Select the date and time for provisioning the service, and click **OK**.

The scheduled date and time are displayed in the MM/DD/YYYY HH:MM AM/PM format.



NOTE:

- You must set a present or future date and time to schedule provisioning of a service. If you select a past time and save the service order, the service is provisioned immediately.
- To delete a provisioning schedule, select the *service-instance-name* on the Service Instances page and click the edit icon. On the Modify *Service-Instance-Name* page, proceed to the Schedule section. In the Schedule section, switch the Schedule toggle off and click **Save**. The provisioning schedule is deleted and the service instance is provisioned immediately.
- Scheduling of service provisioning is a Beta feature in this release.

11. To save and provision the service instance, click one of the following buttons:

- **Save & Provision**—Click this button and confirm when prompted, to save the service instance to the Routing Director database and provision the instance immediately. This option is disabled if you have scheduled the provisioning of a service instance.
- **Save**—Click this button to save the service instance to the Routing Director database. The Service Instances page appears and the new service instance is listed in the Service Instances table. To provision the service instance, select the service instance name and click **Provision**. Alternatively, you can schedule provisioning of the service instance as defined in step 10.



NOTE:

- The **Provision** button is disabled if you schedule the provisioning of a service instance. The service is provisioned automatically at the scheduled date and time.



NOTE: Fields marked with an asterisk (*) are mandatory.

Table 212: Fields in the Add E-Line L2Circuit NSM General Section

Field	Description
Upload JSON File	Click Browse to upload a preconfigured JSON file. You see a message that the file is successfully imported. The values specified in the file are automatically populated in the corresponding UI fields.
Customer*	Enter the name of the customer for whom you are provisioning the service: - If you already added the customer name by using the Customer Inventory page (Orchestration > Service > Customers), click the Customer field to see a drop-down with the customer names that you added. Select the customer name from the drop-down list. - Alternatively, click the Add Customer link to create a new customer. See "Add a Customer" on page 666. The customer name must be unique within an organization. For example, network-operator.
Instance name*	Enter a name for the service instance. For example, l2ckt-171. The instance name can be a set of alphanumeric characters and the special character hyphen (-). The maximum number of characters allowed is 64.

RELATED DOCUMENTATION

[Service Provisioning Workflow | 648](#)

[Add L2 Circuit VPN Nodes | 814](#)

[Modify a Service Instance | 822](#)

Add L2 Circuit VPN Nodes

Use the VPN Nodes section of the Add E-Line L2Circuit NSM page to enter details about the VPN nodes. VPN nodes refer to the PE devices through which the L2 circuit spans in the service provider MPLS network, and which transfer data between the CE devices.

You can configure the following in the VPN Nodes section:

- Details about PE nodes such as the VPN node ID and underlay transport details.
- Site network access details such as the CE device that the PE device connects to, PE interface used to connect to the CE device, connection, and service parameters.
- Signaling type used between PE devices in the service provider network.

Prerequisite

Before adding the L2 circuit node details, you must create an L2 circuit service instance and enter general details about the service. Follow steps 1 to 4 of the ["Add an L2 Circuit Service Instance" on page 810](#) procedure to create a new instance and add L2 circuit service general details.

To add L2 circuit node details in the VPN Nodes section:

1. Click the add (+) icon above the VPN Nodes table.
The VPN Node page appears.
2. Enter the following details in the **General** section:
 - a. Enter the PE device details by referring to the following table:



NOTE:

- You can add only two PE devices for a single-homed L2 circuit. You must add three devices for a multihomed L2 circuit.
- Fields marked with an asterisk (*) are mandatory.

Table 213: PE Device Details (L2 Circuit)

Field	Description
VPN Node ID*	Click the drop-down to select the PE device name for the L2 circuit.
Description	Enter a description about the device, if applicable.

- b. Expand **Underlay Transport** and enter the underlay transport details by referring to the following table:

Table 214: Underlay Transport Details (L2 Circuit)

Field	Description
Transport Instance ID	Enter a unique name for the tunnel or LSP used to transport the L2 circuit. To get the LSP name, log in to the PE router and execute the run <code>show mpls lsp</code> command. The LSP name is displayed in the Ingress LSP > LSPname output field. Enter the LSP name in the Transport Instance ID field. The transport instance ID is an optional field. If you do not provide the LSP name, any available LSP on the PE device is used to send network traffic. NOTE: If you are configuring multihoming on the L2 circuit, do not enter the Transport Instance ID.
Instance Type	Specify the instance type or type of tunnel used between the PE devices in the service provider network to transport the L2 circuit or connection. LSP is the only instance type that is supported in this release.

3. Click **Next** to proceed to the **Network Access** section.
4. Click the add (+) icon above the Network Accesses table.
The Configure Network Access page appears.
5. Configure the network access parameters:
 - a. Enter the name of the CE device and its description by referring to the following table:

Table 215: Network Access Parameters (L2 Circuit)

Field	Description
ID*	Enter a unique ID for the CE device that connects to the PE device interface. For example, ce1.

Table 215: Network Access Parameters (L2 Circuit) (Continued)

Field	Description
Port ID	Select a PE device interface that connects to the CE device from the drop-down. For example, et-0/0/1 NOTE: This drop-down list is empty when you configure network access parameters for the first time. This list is populated when you use the Update Placements option after you exit the VPN Node page. Continue configuring the remaining parameters and you can select the interfaces towards the end of the procedure.
Description	Enter a description about the CE device, if applicable.

- b. Expand **Connection > Encapsulation** and enter the connection parameters by referring to the following table:

Table 216: Connection Parameters (L2 Circuit)

Field	Description
L2 Access Type	Click the L2 Access Type drop-down and select dot1q to specify the type of L2 access used for the L2 circuit in the service provider network. When you select dot1q as the L2 access type, the Update Placements option will display only tagged interfaces for selection in the Port ID drop-down list. If you do not select dot1q as the L2 access type, the Update Placements option will display only untagged interfaces for selection in the Port ID drop-down list. dot1q is the only L2 access type that is supported in this release. The IEEE 802.1Q standard or dot1q is a method of tagging or encapsulating Ethernet frames with VLAN tags. The service provider tags packets between CE and PE devices with one or more customer VLAN (C-VLAN) IDs to identify which VLAN the frame belongs to.

Table 216: Connection Parameters (L2 Circuit) (Continued)

Field	Description
C-VLAN ID	(Optional) Expand Dot1q and in the C-VLAN ID field, enter a unique ID for the C-VLAN. Enter a value from 512 to 4094. NOTE: Enter a C-VLAN ID for the tagged interface that you want to assign for the PE device.

- c. (Optional) Expand **Service > Service Input Bandwidth** and click the add (+) icon above the Input Bandwidth table.
The Input Bandwidth page appears.

- d. Enter the input bandwidth parameters by referring to the following table:

Table 217: Service Input Bandwidth Parameters (L2 Circuit)

Field	Description
Type	Select the type of download bandwidth from the service provider network to the site. Bandwidth per port (bw-per-port) is the only input bandwidth type that is supported in this release. This bandwidth type refers to the input bandwidth per network access (connection) on the CE device.
CBS	Enter the Committed Burst Size (CBS) in bytes per second (Bps). When network traffic does not utilize the configured Committed Information Rate (CIR), the unused bandwidth is accumulated up to a maximum limit defined by the CBS.
CIR	Enter the CIR or the maximum number of bits per second (bps) that an interface can receive. The service provider provides the CIR as the guaranteed bandwidth for network traffic that an interface can receive in one second, under normal line conditions.

- e. Click **OK**.

The input bandwidth entries are listed in the Input bandwidth table.

- f. (Optional) Click the edit or delete icons above the Input Bandwidth table to edit or delete the entries.
- g. (Optional) Expand **Service > Service Output Bandwidth** and click the add (+) icon above the Output Bandwidth table.
The Output Bandwidth page appears.
- h. Enter the output bandwidth parameters by referring to the following table: .

Table 218: Service Output Bandwidth Parameters (L2 Circuit)

Field	Description
Type	Select the type of upload bandwidth from the site to the provider network. Bandwidth per port (bw-per-port) is the only output bandwidth type that is supported in this release. This bandwidth type refers to the output bandwidth per network access (connection) on the CE device.
CBS	Enter the Committed Burst Size (CBS) in bytes per second (Bps). When network traffic does not utilize the configured Committed Information Rate (CIR), the unused bandwidth is accumulated up to a maximum limit defined by the CBS.
CIR	Enter the CIR or the maximum number of bits per second (bps) that an interface can send. The service provider provides the CIR as the guaranteed bandwidth for network traffic that an interface can send in one second, under normal line conditions.

- i. Click **OK**.
The output bandwidth entries are listed in the Output bandwidth table.
 - j. (Optional) Click the edit or delete icons above the Output Bandwidth table to edit or delete the entries.
6. Click **OK** on the Configure Network Access page.
 7. Click **Next** on the VPN Node page to proceed to the Signaling Options section.
 8. Enter the signaling parameters:

- a. Click the **Type** drop-down and click **Ldp-signaling**.

Use this field to define the signaling type that is used between the PE devices in the service provider MPLS network, to create an L2 circuit.

MPLS Label Distribution Protocol (LDP) is the only signaling protocol supported in this release.

- b. Expand **LDP** and click the add (+) icon above the AC and PW Bindings table.
The AC and PW Bindings page appears.

- c. Enter the values by referring to the following table:

Table 219: AC and PW Bindings Parameters (L2 Circuit)

Field	Description
Peer Addr	Enter the IP address of the peer PE device with which the PE device communicates using LDP, to create an L2 circuit through the MPLS network. For example, 10.2.3.4. In a multihomed L2 circuit, enter the IP address of the primary peer.
VC ID	Enter a unique ID for the virtual circuit (VC). The default value of 0 allows placements to auto allocate a label. If you want to explicitly specify a label, enter a number between 1 and 4294967295.

- d. Click **OK**.
The entries are listed in the AC and PW Bindings table.
- e. (Optional) Click the edit or delete icons above the AC and PW Bindings table to edit or delete the entries.
- f. Enable the **Multihoming** toggle on, to enable multihoming on the L2 circuit. In a multihomed L2 circuit there are three devices. For example, consider a multihomed L2 circuit between PE1, PE2, and PE3 devices. PE1 which is connected to CE1 is the local PE access node, and PE2 and PE3, both connected to CE2, are the remote PE peer nodes. The pseudowire between the access node PE1 and the primary peer PE2 is the primary active pseudowire. The pseudowire between access node PE1 and backup peer PE3 is the backup pseudowire. The following two modes are available in a multihomed L2 circuit:
- **Backup mode**—In the default backup mode, there is a single active pseudowire between the access node and the remote PE node. The backup pseudowire is not pre-signalled. If the

active pseudowire fails, then the backup pseudowire becomes active after a connection is established.

- Hot-standby mode—In this mode, the backup pseudowire is pre-signalled. If the active pseudowire fails, then the backup pseudowire becomes active immediately. Since the backup pseudowire is pre-signalled, the switchover time to the backup pseudowire is reduced minimizing traffic disruption.



NOTE: Enter multihoming only on the access node.

Enter multihoming details by referring to the following table:

Table 220: Multihoming

Field	Description
Hot-standby	When disabled, the default backup multihoming mode is activated. When enabled, the hot-standby mode is activated.
Backup	
Peer Addr*	Enter the IP address of the backup peer. For example, 10.10.10.20.
VC Id	Enter a VC label to identify the pseudowire. Enter 0 to allow placements to auto allocate a label. If you want to explicitly specify a label, enter a number between 1 and 4294967295.
Revert timer	
Maximum	Enter time in seconds for maximum period for all traffic to switch over from backup pseudowire to the primary pseudowire after the active pseudowire becomes usable. If traffic has not successfully switched over to the primary pseudowire within the entered time, the switchover is canceled and the backup pseudowire continues to transmit traffic.

Table 220: Multihoming (*Continued*)

Field	Description
Start	Enter time in seconds for the traffic in the backup pseudowire to switch over back to the primary pseudowire, once the primary pseudowire is restored and becomes usable after a failure. For example, if the start time is 10 seconds. The system will wait for 10 seconds after the primary pseudowire becomes usable again, and then start to switch over traffic.

9. Toggle **MTU Allow Mismatch** to True to allow MTU mismatch for a pseudowire. This allows a pseudowire to be brought up between two logical interfaces that are defined on physical interfaces with different MTU values. It is disabled by default.

10. Click **Done** and then click **Close**.

The L2 circuit nodes and site network access parameters you entered are listed in the VPN Nodes table.

11. Configure placement parameters for the service instance.

- a. Click the **Update Placements** button on top of the page and click **Yes** when prompted to save the service instance in its current state.

The service instance is saved to the Routing Director database. The saved instance is listed in the Service Instances table on the Service Instances page.

Service orchestration automatically assigns possible PE device interface placement options for the service instance. You see a message confirming that the placement is updated successfully.

- b. To configure the placement parameters, select the node for which you want to configure placement parameters, and click the edit (pencil) icon.

The VPN Node page appears.

- c. Click **Next** to go to **Network Accesses** page, select the network access ID, and click the edit (pencil) icon.

The Configure Network Access page appears.

- d. Configure the placement parameter by referring to the following table:

Table 221: Placement Parameter

Field	Description
Port ID	Click the drop-down to view the updated interface placement options. Select the PE device interface that connects to the CE device. For example, et-0/0/1. The list displays tagged interfaces when you select dot1q as the L2 access type and untagged interfaces when you do not select the L2 access type.

- e. Click **OK** to exit the Configure Network Access page.

The Network Accesses table displays the interface that you configured. You can also view connection type and C-VLAN ID, if configured.

- f. Click **Next** to view the Signaling Options page and then click **Done** and **Close** to exit the VPN Node page.

The selected interface is displayed in the **Ports** column of the VPN Nodes table.

12. Proceed to step 6 of the "Add an L2 Circuit Service Instance" on page 810 procedure to complete the remaining steps for creating the L2 circuit service instance.

RELATED DOCUMENTATION

[Add an L2 Circuit Service Instance | 810](#)

[Modify a Service Instance | 822](#)

Modify a Service Instance

Prerequisite

- For L3VPN service—"Add an L3VPN Service Instance" on page 773.
- For EVPN and EVPN-VPWS service—"Add an EVPN or EVPN-VPWS Service Instance" on page 791.
- For L2 Circuit service—"Add an L2 Circuit Service Instance" on page 810.

A superuser or network administrator can modify a service instance and provision the modified service in the network. When you provision a modified service instance, Routing Director generates a modify service order to provision the updated service. You can modify a service instance by uploading a preconfigured JSON file containing the changes. Alternatively, you can use the fields in the Routing Director GUI to modify the service-related parameters.



NOTE: You cannot edit the **Customer** and **Instance Name** fields while modifying a service instance.

To modify a service instance:

1. Click **Orchestration > Instances**.
The Service Instances page appears.
2. Select the service instance you want to edit and click the edit (pencil) icon.
The Modify *Service-Instance-Type* page appears.
3. In the General section, click **Browse** to upload a preconfigured JSON file that contains the modifications to the service instance.
Alternatively, follow these steps to modify a service instance by using the GUI fields:
 - For L3VPN, EVPN, and EVPN-VPWS, follow steps 4 and 5.
 - For L2 Circuit, follow steps 4 and 6.
4. In the General section, modify the values in the fields, if applicable.
5. (For L3VPN, EVPN, and EVPN-VPWS) Modify the site settings:
 - a. Click **Next** to proceed to the Site Settings section.
 - b. In the Site Settings section, select the site that you want to modify and click the edit (pencil) icon.
The Edit Site page appears.
 - c. Modify the site details:
 - Edit L3VPN service site details by referring to [Table 186 on page 778](#).
 - Edit EVPN and EVPN-VPWS service site details by referring to [Table 199 on page 797](#).
 - d. Click **OK**.
To view the modified site details, expand the site name and expand **Properties**.
 - e. To edit the site network access details, expand the site name and expand **Site Network Access**.
 - f. Select the site network access that you want to modify and click the edit (pencil) icon.
The Edit Connection page appears.
 - g. Expand each section on the page to edit the field values as applicable:

- Edit L3VPN site network access details by referring to ["Add L3VPN Site Network Access Parameters" on page 779](#).
 - Edit EVPN and EVPN-VPWS site network access details by referring to ["Add EVPN or EVPN-VPWS Site Network Access Parameters" on page 799](#).
- h. Modify the placement configurations by using one of the following options:
- **Placement** section is displayed on the Edit Connection page—Expand **Placement** and modify the placement parameters by clicking the drop-down in each field. [Table 222 on page 825](#) lists the placement parameters.
 - **Placement** section is not displayed on the Edit Connection page:
 - i. Click the **Update Placements** button on top of the page and click **Yes** when prompted to confirm.

The service instance is saved to the Routing Director database and service orchestration automatically assigns possible placement options such as PE devices and interfaces, CE devices, and VLANs for the service instance. You see a message that the placement was successfully updated.
 - ii. To modify the placement configurations, select the *service-instance-name* on the Service Instances page and click the edit icon.

The Modify *Service-Instance-Name* page appears.
 - iii. Proceed to the Site Settings section, select the site network access that you want to modify and click the edit icon.

The Edit Connection page appears and the **Placement** section is displayed on the page.
 - iv. Expand **Placement** and modify the placement parameters by referring to [Table 222 on page 825](#) for L3VPN services or by referring to [Table 223 on page 825](#) for EVPN and EVPN-VPWS services.



NOTE: The **Placement** section is not available if:

- You modify site details such as site location and postal code.
- Service orchestration could not allocate the PE device as a resource for the VPN service.
- The execution state of the service order generated for the service is **Not available**.

For L3VPN services:

Table 222: Placement Parameters for L3VPN Services

Field	Description
CE Node	Click the CE Node drop-down and select the CE device that you want to connect to the PE device for the VPN service.
PE Node	Click the PE Node drop-down and select the PE device that you want to allocate for the VPN service.
PE Interface	Click the PE Interface drop-down and select the PE interface name that you want to assign for the VPN service.

For EVPN and EVPN-VPWS services:

Table 223: Placement Parameters for EVPN and EVPN-VPWS Services

Field	Description
CE Node	Click the CE Node drop-down and select the CE device that you want to connect to the PE device for the VPN service.
PE Node	Click the PE Node drop-down and select the PE device that you want to allocate for the VPN service.
PE Interface	Click the PE Interface drop-down and select the PE interface name that you want to assign for the VPN service.

- i. (Optional) Click **Reset Placements** to remove the configured placement parameters at the site network access level and delete the **Placements** section from the page.
 - j. Click **OK** to save the modified site network access details. The modified values are listed in the Site Network Access table.
Alternatively, click **Cancel** to exit the Edit Connection page without saving the changes you made.
6. (For L2 Circuit) Modify the placement configurations and VPN node settings:
- a. Click **Next** to proceed to the VPN Nodes section.

- b. Click the **Update Placements** button on top of the page and click **Yes** when prompted to save the service instance in its current state.
- c. Select the node for which you want to configure placement parameters, and click the edit (pencil) icon.

The VPN Node page appears.

- d. Modify placement configuration and VPN node settings for the service instance by following Step 2 through Step 10 of the ["Add L2 Circuit VPN Nodes" on page 814](#) procedure.



NOTE: You cannot edit the VPN node ID while modifying the node details.

7. Click **Next** to proceed to the Summary section.

You see a graphical map of your service topology and a summary of the service order.

8. (Optional) Click **Export** to export and save the service order in JSON format on your local system.
9. (Optional) Click **Reset Placements** to remove existing placement configurations for the service instance.

Provisioning of a service instance might fail when existing placement configurations for the service instance are no longer available or applicable. When you reset placement parameters, service orchestration automatically assigns placement options for the instance before provisioning the service.

10. Click **Next** to proceed to the Schedule section.

11. (Optional) Switch the Schedule toggle on if you want to schedule when to provision the modified service. The Schedule option is disabled by default. When you schedule provisioning, the **Save & Provision** button on top of the page is disabled and grayed out.

The Schedule field appears. Perform the following steps to schedule the date and time to provision the service:

- a. Click the calendar icon in the Schedule field.
- b. Select the date and time for provisioning the service, and click **OK**.

The scheduled date and time are displayed in the MM/DD/YYYY HH:MM AM/PM format.



NOTE:

- You must set a present or future date and time to schedule provisioning of a service. If you select a past time and save the service order, the service is provisioned immediately.

- To delete a provisioning schedule, select the *service-instance-name* on the Service Instances page and click the edit icon. On the Modify *Service-Instance-Name* page, proceed to the Schedule section. In the Schedule section, switch the Schedule toggle off and click **Save**. The provisioning schedule is deleted and the service instance is provisioned immediately.
- Scheduling of service provisioning is a Beta feature in this release.

12. To save and provision the service instance, click one of the following buttons:

- **Save & Provision**—Click this button and confirm when prompted, to save the service instance to the Routing Director database and provision the instance immediately. This option is disabled if you have scheduled the provisioning of a service instance.
- **Save**—Click this button to save the service instance to the Routing Director database. The Service Instances page appears and the modified service instance is listed in the Service Instances table.

To provision the service instance automatically, configure the provisioning schedule as defined in step 11. To manually provision the service, select the service instance name in the Service Instances table and click the **Provision** button.



NOTE: The **Provision** button is disabled if you schedule the provisioning of a service instance. The service is provisioned automatically at the scheduled date and time.

RELATED DOCUMENTATION

[Service Provisioning Workflow | 648](#)

[Add an L3VPN Service Instance | 773](#)

[Add L3VPN Service Site Details | 777](#)

[Add an EVPN or EVPN-VPWS Service Instance | 791](#)

[Add EVPN or EVPN-VPWS Service Site Details | 796](#)

[Add an L2 Circuit Service Instance | 810](#)

[Add L2 Circuit VPN Nodes | 814](#)

Manage Service Instances

IN THIS CHAPTER

- [Service Instance Overview | 828](#)
- [About the Service Instances Page | 829](#)
- [View Service Instance Details | 836](#)
- [View Passive Assurance Monitoring Data | 843](#)
- [View EVPN or EVPN-VPWS Passive Assurance Monitoring Results | 845](#)
- [View L2 Circuit Passive Assurance Monitoring Data | 852](#)
- [View L3VPN Passive Assurance Monitoring Results | 854](#)
- [Physical Interfaces Health Check Results | 864](#)
- [Logical Interfaces Health Check Results | 869](#)
- [Performance Graphs to View Passive Assurance Monitoring Data | 871](#)

Service Instance Overview

A service instance is a detailed description about a service. Typically written in the JSON or YAML format, a service instance defines the elements of a service. For example, an L3VPN service instance defines elements such as the customer name, service topology, and the sites where the L3VPN service is to be deployed. You create service instances for onboarding devices and defining network services, based on the corresponding service designs that Routing Director provides.

A service order is an instruction to provision a service instance in the network. Routing Director generates a service order when you create or modify a service instance and provision the instance. A service order is also generated for deleting or deprovisioning a service instance. Each type of service order (create, modify, and delete) has an automated workflow associated with it. A service instance can have many service orders associated with it.

When you create or modify a service instance and save it, the instance is uploaded to the Routing Director database. The status of the service instance is *Uploaded*. To generate a service order for the instance and activate the automated workflow associated with the service order, you must provision the service instance. Routing Director automatically provisions and monitors the service after a provisioning

workflow is successfully executed. The workflow includes a series of tasks such as transforming the service order, committing device configuration, and so on, for provisioning the service.

The Service Instances page (**Orchestration > Service > Instances**) lists all the service instances created in an organization.

RELATED DOCUMENTATION

[About the Service Instances Page | 829](#)

About the Service Instances Page

IN THIS SECTION

- [Tasks You Can Perform | 829](#)
- [Field Descriptions | 832](#)

The Service Instances page lists the services that you (a service provider) have created in an organization and intend to provision for your customers. The services include service instances for L3VPN, EVPN, EVPN-VPWS, or L2 circuit service provisioning. The service instances are listed by the instance name, customer name, the associated service design, devices on which the service is provisioned, and so on.

To access the **Service Instances** page, click **Orchestration > Instances** on the navigation menu.

Tasks You Can Perform

You can perform the following tasks on the **Service Instances** page:

- View details of service instances.

Use one of the following options to view details about a service instance:

- To view service details in a pop-up pane on the right:
 - Hover over the service instance and click the **Details** icon that appears next to the service instance name.
 - Select the service instance and click **More > Detail**.

The *service-instance-name* pane appears on the right side of the GUI. Use the scroll bar at the bottom to scroll left to right or click **X** to close the pane. For information about the details of the fields, see [Table 225 on page 835](#).

To view the service instance details in JSON format, click the **JSON View** hyperlink on the details pane.

To view the network resources allocated for the service instance, click the **Placement Resources** hyperlink on the details pane.

- Click the *service-instance-name* hyperlink and view details on the *service-instance-name* Details page. See ["View Service Instance Details" on page 836](#).
- Create service instances:
 - Create L3VPN service instances. See ["Add an L3VPN Service Instance" on page 773](#).
 - Create EVPN or EVPN-VPWS service instances. See ["Add an EVPN or EVPN-VPWS Service Instance" on page 791](#).
 - Create L2 circuit service instances. See ["Add an L2 Circuit Service Instance" on page 810](#).
- Modify service instances. See ["Modify a Service Instance" on page 822](#).

- Provision a service instance

To provision a service instance, select the service instance and click **Provision**. Routing Director generates a corresponding service order to activate the provisioning workflow. You can view the service order in the Service Orders page (**Orchestration > Monitoring > Service Orders**) and monitor the execution of the corresponding workflow on the Workflows (**Orchestration > Monitoring > Workflows**) page.



NOTE: You can also schedule the date and time for provisioning a service when you are creating or modifying a service instance. The **Provision** option is disabled if you schedule the provisioning of a service instance. The service is provisioned automatically at the scheduled date and time.

- Deprovision a service—To delete or deprovision a service from your network, select the service instance and click **More > Deprovision**. Click **Yes** when prompted for confirmation.
- Force Sync a service configuration.

To use the Force Sync option, click **More > Force Sync**.

The **Force Sync** option reverts a service configuration to the configuration defined by the service instance used to provision the service. Use this option when you have made changes to the service configuration and want to revert to the original configuration defined by the service instance.

- Upgrade the version of the service design that is associated with a service instance. See ["Upgrade Service Design Version for a Service Instance" on page 658](#).
- Export details of a service instance in the JSON format—To export a service instance in the JSON format and save it on your local system, select the service instance and click **More > Export**. The `plan.service-instance-name.json` file is downloaded to your local system.
- Assign placement options for service instances—To automatically assign placement or resource allocation options for an L3VPN, EVPN, or EVPN-VPWS service instance, select the service instance and click **More > Update Placements**.

Placement options can also be assigned when you are creating or modifying a service instance by using the GUI or service orchestration cMGD CLI. See ["Manage Placement Configurations for Service Instances" on page 768](#).

- Add a tag—A tag is a label in the key:value format that you add to a service instance so that you can group service instances. For example, you can assign the tag customer:ABC to all service instances associated with customer ABC.

To add one or more tags to a service instance:

1. Click **More > Add Tags**.

The Manage Tags page appears.

2. In the Tags field, select a tag to be applied from the drop-down list.

Alternatively, you can also create a new tag by clicking the Create New Tag link at the bottom of the page and apply the new tag. If you are creating a new tag, ensure that:

- The key starts with a letter and is followed by letters, numbers, periods (.), hyphens (-), or underscores (_), and is between 1 and 200 characters long.
- The value includes alphanumeric characters, periods (.), hyphens (-), and underscores (_), and is between 1 and 200 characters long.

You can view all the tags applied to a service instance in the Tags column of the Service Instances page.

- Filter the data displayed in the table—Click the filter icon (funnel) and select whether you want to show or hide advanced filters. You can then add or remove filter criteria, save criteria as a filter, apply or clear filters, and so on. The filtered results are displayed on the same page.
- Search by using keywords—Click the search icon (magnifying glass), enter the search term in the text box, and press Enter. The search results are displayed on the same page.
- Show or hide columns in the table or reset page preferences, using the vertical ellipsis menu.
- Sort, resize, or re-arrange columns in a table (grid).



NOTE: Sort functionality is not available for the Customer, Devices, Device Count, Sites, and Last Modified By columns.

- Set the number of rows to be displayed in the table—Click the **Display** drop-down at the bottom-right corner of the Service Instances table to set the number of rows (15, 30, 50, or 100) to be displayed in the table. The default number of rows displayed is 15.
- Set auto refresh frequency—Click the Refresh icon at the bottom-left corner of the Service Instances table to view the most recent updates.

Use the arrow heads to traverse through the pages or enter a page number to view the entries listed on that page.

You can also auto refresh the table to view updated details. To set auto refresh frequency:

1. Click the horizontal ellipsis near the Refresh icon.
2. Select auto refresh frequency from the available options (30 seconds, 1 minute, 3 minutes, and 5 minutes).

The next refresh time is displayed near the Refresh icon.

To turn off auto refresh, click the horizontal ellipsis and click **Off**.

Field Descriptions

Table 224 on page 832 lists the fields on the Service Instances page.

Table 224: Fields on the Service Instances Page

Field	Description
Name	Name of the service instance.
Customer Id	Unique ID of the customer for whom you are provisioning the service.
Customer	Name of the customer for whom you are provisioning the service.
Service Design	Name of the service design associated with the service instance.

Table 224: Fields on the Service Instances Page (*Continued*)

Field	Description
Design version	Service design version associated with the service instance.
Devices	Hostnames of the devices on which you are provisioning the service.
Device Count	Number of devices on which you are provisioning the service.
Customer Sites	Customer sites that connect to the service.
Order State	The execution state of the service order generated for the service instance. A service order can have the following execution states: • Processing—The service order is in the process of being executed. • Success—The service order is successfully executed. • Success with warnings—The service order is successfully executed but certain tasks in the workflow have failed. You can view the tasks that have failed and take appropriate action if required. To view detailed task logs, navigate to the Orchestration > Workflows > Workflow-Run-Name page. See "View Workflow Run Details" on page 885. • Failed—The service order has failed to execute. For example, a service order fails to execute if a PE device through which the service spans is disconnected. The order state is then displayed as Failed.

Table 224: Fields on the Service Instances Page (*Continued*)

Field	Description
Status	The provisioning status of the service instance. A service instance can have one of the following provisioning states: • Feasibility Checked—The service instance is created, but the provisioning workflow has not been activated. • Scheduled—The provisioning workflow is scheduled to be executed at a future date and time. You must set a present or future date and time to schedule provisioning of a service. When a service instance is scheduled, it is locked and cannot be modified until the schedule is canceled or executed. • Designed—Possible placement or resource allocation options are assigned for the service instance, but the resources are not yet reserved or allocated. • Reserved—Resource allocation for the service instance is complete. • Provisioning—The service instance is being provisioned with the allocated resources. • Active—The service configurations are successfully applied on the devices. • Inactive—The service is deactivated and only the delete service order is accepted for the deactivated service. • Blocked—Service configurations failed to apply on any device or on Active Assurance and Insights instances. • Terminated—The service instance has been deleted.
Status Message	Additional information about the status of the service order initiated for the service instance. The status is Uploaded after the service instance is saved (either after creating or modifying).

Table 224: Fields on the Service Instances Page (*Continued*)

Field	Description
Scheduled Time	The date and time when the service instance is scheduled to be provisioned. The date and time are displayed only if you have scheduled provisioning of the service instance.
Last Modified	The date and time when the service instance was last modified.
Tags	Tags assigned to the service instance. If there are more than one tag assigned to the service instance, an integer is present next to the displayed tag. For example, +2 in the Tags column indicates that there are two more tags in addition to what is currently listed under Tags. You can hover over the integer to view the additional tags assigned to the service instance.
Last Modified By	Name of the user who last modified the service instance.

Table 225: Fields on the *service-instance-name* Pane

Field	Description
Name	Name of the service instance.
Customer	Name of the customer for whom you are provisioning the service.
Service Design	Name of the service design associated with the service instance.
Design Version	Service design version associated with the service instance.
Customer Sites	Customer sites that connect to the service.

Table 225: Fields on the *service-instance-name* Pane (Continued)

Field	Description
Status Message	Information about the status of the service order initiated for the service instance. The status is Uploaded after the service instance is saved (either after creating or modifying).
Order ID	Unique alphanumeric ID assigned to the service order generated for the service instance.
Scheduled Time	The date and time when the service instance is scheduled to be provisioned. The date and time are displayed only if you have scheduled provisioning of the service instance.
Last Modified	The date and time when the service instance was last modified.
Last Modified By	Name of the user who last modified the service instance.

RELATED DOCUMENTATION

[Add Network Resources for Provisioning Services | 670](#)

[Add an L3VPN Service Instance | 773](#)

[Add an EVPN or EVPN-VPWS Service Instance | 791](#)

[Add an L2 Circuit Service Instance | 810](#)

[Modify a Service Instance | 822](#)

View Service Instance Details

IN THIS SECTION

[About the Active Assurance Tab | 839](#)

- [About the Order History Tab | 840](#)
- [About the Passive Assurance tab | 841](#)
- [About the Configurations Tab | 842](#)

The *Service-Instance-Name* Details page displays details about service instances that you have created.

To navigate to the *Service-Instance-Name* Details page:

1. Click **Orchestration > Instances**.

The Service Instances page appears.

2. Select a service instance and click the *service-instance-name* hyperlink.

The *Service-Instance-Name* Details page appears.

You can view service instance details in the following tabs on the *Service-Instance-Name* Details page:

- **Properties**—View general details about the service instance such as the name of the service instance, associated service design, customer name, and so on. See [Table 226 on page 838](#).
- **Order Status**—View status of the last service order initiated for the instance. Expand **Components** and **Workflows** to view the execution details of the associated service order by the service components (placement, device, insights, and active assurance) and order workflow runs. You can view:
 - The date and time when the tasks related to the service component was executed.
 - The date, time, number of times the tasks in the service order workflow was run, and the status of the task execution.

Click the **View Workflow Run** hyperlink to view detailed workflow task logs on the Workflow Runs List page. See ["View Workflow Run Details" on page 885](#).

- **Active Assurance**—View active assurance Tests and Monitors data to monitor the quality of L3VPN services after provisioning, under the Active Assurance tab. See ["About the Active Assurance Tab" on page 839](#).



NOTE: The Active Assurance tab is displayed for L3VPN service instances only.

- **Order History**—View history of all service orders generated for the service instance. See ["About the Order History Tab" on page 840](#).

- **Passive Assurance**—View passive assurance monitoring data to monitor the health of VPN services after provisioning, under the Passive Assurance tab. See ["About the Passive Assurance tab" on page 841](#).
- **Configurations**—View configurations related to devices on which the service instance is provisioned and service monitoring. See ["About the Configurations Tab" on page 842](#).
- **Map**—View the map of the service topology that is defined in the service instance.

[Table 226 on page 838](#) lists the service instance properties displayed on the *service-instance* Details page.

Table 226: Fields in the Properties Tab

Field	Description
Name	Name of the service instance and customer for whom the service is provisioned. The name is displayed in the <i>service-instance:customer-name</i> format.
Customer	Name of the customer for whom you are provisioning a service.
Service Design	Name of the service design associated with the service instance.
Design Version	Service design version associated with the service instance.
Sites	Sites through which the service spans.
Status Message	Status of the service order initiated for the service instance. For example, for a service order that is uploaded but not yet provisioned, the status displays as Uploaded. For a device onboarding service order that is executed, the status displays as Applied configuration to <i>device-name</i>.
Order ID	The ID of the last executed service order for the service instance.

Table 226: Fields in the Properties Tab *(Continued)*

Field	Description
Scheduled Time	The date and time when the service instance is scheduled to be provisioned. Scheduling of service instance provisioning is optional. This field does not display any entry if the service instance provisioning is not scheduled.
Last Modified	The date and time when the service instance was last modified.
Last Modified By	The user who last modified the service instance.

About the Active Assurance Tab

The Active Assurance tab provides information about the status of Active Assurance Tests and Monitors that are used to monitor the service after it is provisioned, by measuring network traffic for the service.

See [Table 227 on page 839](#) for information about the fields.



NOTE: The Active Assurance tab is displayed for L3VPN service instances only.

Table 227: Fields in the Active Assurance Tab

Field	Description
Tests	The Tests column displays the following data: • Total number of Tests run. • Number of Tests that have passed. • Number of Tests that have failed. • Number of Tests that could not run due to an error. To view details about the Tests, click the hyperlinked number for total number of Tests run. You are directed to the Tests page (Observability > Active Assurance > Tests) where you can see details of the Tests. See "About the Tests Page" on page 1100.

Table 227: Fields in the Active Assurance Tab (Continued)

Field	Description
Monitors	The Monitors column displays the following data: • Total number of Monitors. • Number of Monitors running. • Number of Monitors that have stopped running. • Number of Monitors that failed due to errors. To view details about the Monitors, click the hyperlinked number for total number of Monitors. You are directed to the Monitors page (Observability > Active Assurance > Monitors) where you can see details of the Monitors. See "About the Monitors Page" on page 1247.

About the Order History Tab

The Order History tab displays a summary of all service orders associated with the service instance. You can view the date and time when an order was initiated and additional details about each service order, as described in [Table 228 on page 840](#).

Table 228: Fields in the Order History Tab

Field	Description
Operation	The type of operation specified in the order (create, modify, or delete).

Table 228: Fields in the Order History Tab (*Continued*)

Field	Description
State	State of execution of the service order. A service order can have the following execution states: • Processing—The service order is in the process of being executed. • Success—The service order is successfully executed. • Success with warnings—The service order is successfully executed but certain tasks in the workflow have failed. You can view the logs of tasks that fail and take further action if required. For information about viewing task logs, see "View Workflow Run Details" on page 885. • Failed—The service order has failed to execute.
Status	Status of the service order. For example, for a service order that is uploaded but not yet provisioned, the status is Uploaded. For a device onboarding service order that is executed, the status displays as Applied configuration to <i>device-name</i>.
Device Count	Number of devices on which the service order is executed.
Edited By	Name of the user who initiated the order.
SD Version	Service design version associated with the service order.
View Content	You can view the contents of the service order by clicking the View Content hyperlink.

About the Passive Assurance tab

The Passive Assurance tab displays the results of health checks that Routing Director runs after a VPN service is provisioned. Routing Director monitors the health of the service and its components, such as physical and logical interfaces. For more information about the data displayed in the Passive Assurance tab, see ["View Passive Assurance Monitoring Data"](#) on page 843.

About the Configurations Tab

The Configurations tab displays the configurations for the service components (device, Insights, and Active Assurance). The service components appear on the left side of the tab. Click a component to view the component-specific configuration on the right side of the tab. See [Table 229 on page 842](#) for information about the fields in the Configurations tab.

The right pane also displays the following tabs:

- Outline—Displays the summarized outline of the selected configuration file. Expand the outline sections to view more details.
- Search—Allows you to search for specific content in the configurations file. Enter the search term in the text field and press Enter. The search results are displayed in the same section.

Table 229: Fields in the Configurations Tab

Field	Description
Device	Displays the list of devices on which the service is provisioned. Click a <i>device-name</i> to view the device configurations on the right side of the tab.
Insights	Displays the passive monitoring configurations for monitoring the service. Click the <i>service-order</i> to view the configurations on the right pane.
Active Assurance	Displays the Active Assurance configurations for testing the service. Click the <i>service-order</i> to view the configurations on the right pane.

RELATED DOCUMENTATION

About the Service Instances Page 829
About the Tests Page 1100
About the Monitors Page 1247
View Workflow Run Details 885
View Passive Assurance Monitoring Data 843

View Passive Assurance Monitoring Data

IN THIS SECTION

- [Tasks You Can Perform in the Passive Assurance Tab | 843](#)

Routing Director automatically monitors the health of a VPN service after provisioning the service in the network. To view the health check results for a VPN service and its components:

1. Click **Orchestration > Instances**.

The Service Instances page appears.

2. Select the service instance and click the *service-instance-name* hyperlink.

The *Service-Instance-Name* Details page appears.

3. Click the **Passive Assurance** tab.

The monitoring data of the service components (categorized into accordions) generated by health checks appear.

You can view passive assurance monitoring data for the following VPN services and related components:

- EVPN service. See ["View EVPN or EVPN-VPWS Passive Assurance Monitoring Results" on page 845](#).
- L2 circuit service. See ["View L2 Circuit Passive Assurance Monitoring Data" on page 852](#).
- L3VPN service. See ["View L3VPN Passive Assurance Monitoring Results" on page 854](#).

Tasks You Can Perform in the Passive Assurance Tab

You can perform the following tasks in the Passive Assurance tab:

- View service component parameters—Expand the accordions to view the parameters for the corresponding service component. The parameters are displayed in a tabular format. Table entries for parameters for which you can view detailed health check results are hyperlinked. Click any hyperlinked entry to view the corresponding performance graph and events (alerts) generated for the parameter in the *service-component-name* Details for *service-instance-name* page.
- Set the number of rows to be displayed in the accordion tables.

Click the **Display** drop-down at the bottom-right corner of the accordion tables to set the number of rows (15, 30, 50, or 100) to be displayed in the table. The default number of rows displayed is 15.

- Refresh data displayed in the accordion tables—Click the Refresh icon at the bottom-left corner of the accordion tables to view the most recent updates.
- Show or hide columns in accordion tables or reset page preferences, using the vertical ellipsis menu.
- View events (alerts) generated for service components.

The Relevant Events section on the right side of the accordion displays the last two events (alerts) in order of their severity. To view details about the events, do any one of the following:

- If there are fewer than two events, hover over the **View Details** hyperlink. You see a pop-up with the device name, description, created time, last received time, and number of recurrences of the event.
- If there are more than two events, click the **View all Relevant Events** hyperlink to view the complete list of events with details such as severity, timestamp, type, recurrence, and description of each event for the accordion category. You can view relevant events for the past seven days.

The severity level of events is defined by the following categories.

- Urgent Action Needed (red icon)—Indicates that a critical event has occurred. The functioning of the component is affected and needs urgent user intervention to fix the issue.
- Action Needed (orange icon)—Indicates that a major event has occurred on the component and needs user action but not urgently. The functioning of the component is affected but not drastically.
- Being Monitored (yellow icon)—Indicates that a minor event has occurred on the component but needs no user action. The component is being monitored.

When no event is reported and the component is functioning normally, the status is displayed as **No alerts** (green icon).

RELATED DOCUMENTATION

[Physical Interfaces Health Check Results | 864](#)

[Logical Interfaces Health Check Results | 869](#)

[View L2 Circuit Passive Assurance Monitoring Data | 852](#)

[View Service Instance Details | 836](#)

View EVPN or EVPN-VPWS Passive Assurance Monitoring Results

IN THIS SECTION

- [View L2VPN or EVPN-VPWS Health Check Results | 846](#)
- [View OAM Health Check Results | 848](#)
- [View LLDP Health Check Results | 850](#)

After provisioning an EVPN or EVPN-VPWS service instance in the network, Routing Director automatically monitors the health of the service and its related parameters. You can view the health check results and the alerts generated to indicate if any user action is required. The monitoring data is categorized into accordions and each accordion displays related information and alerts generated.

You can view the following accordions for EVPN or EVPN-VPWS service instances:

- Physical Interfaces—View health check results for physical interfaces of PE devices through which the VPN service spans. Health check results include parameters such as:
 - Link flaps
 - Link state
 - Input and output bandwidth utilization
 - Input and output error counts
 - Interface speed

See ["Physical Interfaces Health Check Results" on page 864](#).

- Logical Interfaces—View health check results for operation status of the logical interfaces on the PE devices through which the VPN service spans. See ["Logical Interfaces Health Check Results" on page 869](#).
- L2VPN or EVPN-VPWS—View health check results for L2VPN or EVPN-VPWS services. See ["View L2VPN or EVPN-VPWS Health Check Results" on page 846](#).
- OAM— View health check results for the OAM accordion. See ["View OAM Health Check Results" on page 848](#).
- LLDP—View health check results for the LLDP accordion. See ["View LLDP Health Check Results" on page 850](#).



NOTE: KPI-related data may not be displayed in the accordion tables and on the graphs when:

- There are device errors
- The data displayed is derived from or is dependent on another KPI which is not being streamed by the devices

View L2VPN or EVPN-VPWS Health Check Results

IN THIS SECTION

- [View Health Check Results for L2VPN or EVPN-VPWS Services | 847](#)

The L2VPN or EVPN-VPWS accordion displays information about L2VPN or EVPN-VPWS service related components.

To access the L2VPN or EVPN-VPWS accordion:

1. Navigate to the **Orchestration > Instances > Service Instances** page.
2. Select a service instance and click the *service-instance-name* hyperlink to access the *Service-Instance-Name* Details page.
3. Click the Passive Assurance tab.
4. Click to expand the L2VPN or EVPN-VPWS accordion.

[Table 1 on page 846](#) lists the fields in the L2VPN or EVPN-VPWS table.

To view health check results for L2VPN or EVPN-VPWS services, click any hyperlinked entry in the table to open the corresponding accordion in the L2VPN or EVPN-VPWS Detail page.

Table 230: Fields in the L2VPN or EVPN-VPWS Table

Field	Description
VPN Network Access ID	Unique ID for the connection between the PE device and the CE device.

Table 230: Fields in the L2VPN or EVPN-VPWS Table (Continued)

Field	Description
Neighbor Count	Number of neighbors on PE device with which the PE device connects to create the L2VPN service in the network.
Device Name	Name of the PE device on which the L2VPN or EVPN-VPWS service is provisioned.
Interface Name	Name of the interface of the PE device connected to the CE device. A logical interface is created on a physical interface of the device.
Instance IFL No	Logical interface instance of a PE device; for example, 100.
Interface Status	Status of the logical interface of a PE device.
Routing Instance Name	Name of the routing instance configured on the PE device.

View Health Check Results for L2VPN or EVPN-VPWS Services

The L2VPN or EVPN-VPWS Detail page displays the health check results for the neighbor count and interface status of L2VPN or EVPN-VPWS services. You can view the name of the PE device and the name of the interface on the PE device at the top of the page. You can also view IP address of the CE device, instance name, and IP address of the interface that connects the PE device to a CE device. The accordion heading displays the parameter name, alerts generated for the parameter, and a condensed view of the performance graph displaying the health and functioning of the parameter.

To view the health check results in detail, click to expand the corresponding accordion. The accordion displays the performance graph for the neighbor count and interface status of L2VPN or EVPN-VPWS services. For more information about performance graphs, see ["Performance Graphs to View Passive Assurance Monitoring Data" on page 871](#).

[Table 2 on page 848](#) lists the accordions on the L2VPN or EVPN-VPWS Detail page.

Table 231: Accordions on the L2VPN or EVPN-VPWS Detail Page

Accordion	Description
Neighbor Count	View information and events related to the number of EVPN neighbors on the interface that connects the PE device to another PE device(s). The x-axis (horizontal) on the performance graph displays the time range for which information is displayed. The y-axis (vertical) displays the number of neighbors. The graph displays a line with the data points related to events that have occurred.
Interface Status	View information and events related to the status of the interface from a PE device to a CE device. The x-axis (horizontal) on the performance graph displays the time range for which information is displayed. The y-axis (vertical) displays the status of the interface. The graph displays a line with the data points related to events that have occurred.

View OAM Health Check Results

IN THIS SECTION

- [View Health Check Results for OAM Parameters | 849](#)

The operation, administration, and management (OAM) accordion displays information about link fault management (LFM) discovery state of the OAM interface.

To access the OAM accordion:

1. Navigate to the **Orchestration > Instances > Service Instances** page.
2. Select a service instance and click the *service-instance-name* hyperlink to access the *Service-Instance-Name* Details page.
3. Click the Passive Assurance tab.
4. Click to expand the OAM accordion.

[Table 3 on page 849](#) lists the field in the OAM accordion table.

To view health check results for the OAM accordion, click any hyperlinked entry in the OAM accordion table to open the corresponding accordion in the OAM Details page.

Table 232: Fields in the OAM Table

Field	Description
Device Name	Name of the PE device on which the service is provisioned.
LFM Interface Name	Name of the LFM interface on which OAM is enabled.
LFM Discovery State	Discovery states of the LFM interfaces. The discovery states: • Send Any—The LFM is configured between PE and CE interfaces. • Active—The interface discovers and monitors the peer on the link. • Passive—The peer initiates the discovery process.

View Health Check Results for OAM Parameters

The OAM Details page displays the health check results for the LFM discovery state of the OAM interface, categorized into accordion. The accordion heading displays the parameter name, alert generated for the parameter, and a condensed view of the performance graph displaying the health and functioning of the parameter.

To view the health check results in detail, click to expand the accordion. The accordion displays the performance graph for the discovery state of the LFM interfaces. For more information about performance graphs, see ["Performance Graphs to View Passive Assurance Monitoring Data" on page 871](#).

[Table 4 on page 850](#) lists the accordion on the OAM Details page.

Table 233: Accordion on the OAM Detail Page

Accordion	Description
LFM Discovery State	View information and events related to the discovery state of the LFM interfaces. The x-axis (horizontal) on the performance graph displays the time range for which information is displayed. The y-axis (vertical) displays the LFM discovery state. The graph displays a line with the data points related to events that have occurred.

View LLDP Health Check Results

IN THIS SECTION

- [View Health Check Results for LLDP Parameters | 851](#)

The Link Layer Discovery Protocol (LLDP) accordion displays information about LLDP status in the EVPN service.

To access the LLDP accordion:

1. Navigate to the **Orchestration > Instances > Service Instances** page.
2. Select a service instance and click the *service-instance-name* hyperlink to access the *Service-Instance-Name* Details page.
3. Click the Passive Assurance tab.
4. Click to expand the LLDP accordion.

[Table 5 on page 851](#) lists the fields in the LLDP table.

To view health check results for the LLDP status, click any hyperlinked entry in the LLDP table to open the corresponding accordion in the LLDP Detail page.

Table 234: Fields in the LLDP Table

Field	Description
Device Name	Name of the PE device on which the service is provisioned.
Interface Name	Name of the interface of PE device.
LLDP Enabled	LLDP enabled status of the interface on the PE device.

View Health Check Results for LLDP Parameters

The LLDP Details page displays the health check results for the LLDP status in the EVPN service, categorized into accordion. The accordion heading displays the parameter name, alert generated for the parameter, and a condensed view of the performance graph displaying the health and functioning of the parameter.

To view the health check results in detail, click to expand the accordion. The accordion displays the performance graph for the LLDP enabled status on the interface. For more information about performance graphs, see ["Performance Graphs to View Passive Assurance Monitoring Data"](#) on page 871.

[Table 6 on page 851](#) lists the accordion on the OAM Details page.

Table 235: Accordions on the LLDP Details Page

Accordion	Description
LLDP Enabled	View information and events related to the LLDP enabled status of the interface on the PE device. The x-axis (horizontal) on the performance graph displays the time range for which information is displayed. The y-axis (vertical) displays the LLDP enabled interface. The graph displays a line with the data points related to events that have occurred.

View L2 Circuit Passive Assurance Monitoring Data

IN THIS SECTION

- [View L2 Circuit Service Health Check Results | 852](#)
- [View Health Check Results for L2 Circuit Connection Status | 854](#)

After provisioning an L2 circuit service instance in the network, Routing Director automatically monitors the health of the service and its related components. You can view the health check results and the alerts generated to indicate if any user action is required. The monitoring data is categorized into accordions and each accordion displays related information and alerts generated.

You can view the following accordions for L2 circuit service instances:

- **Physical Interfaces**—View health check results for physical interfaces of PE devices through which the VPN service spans. Health check results include parameters such as:
 - Link flaps
 - Link state
 - Input and output bandwidth utilization
 - Input and output error counts
 - Interface speed

See ["Physical Interfaces Health Check Results" on page 864](#).

- **Logical Interfaces**—View health check results for operation status of the logical interfaces through which the VPN service spans. See ["Logical Interfaces Health Check Results" on page 869](#).
- **L2 Circuit**—View health check results for the L2 circuit connection status. See ["View L2 Circuit Service Health Check Results" on page 852](#).

View L2 Circuit Service Health Check Results

The L2 Circuit accordion displays the L2 circuit service related components.

To access the L2 Circuit accordion:

1. Navigate to the **Orchestration > Instances > Service Instances** page.

2. Select a service instance and click the *service-instance-name* hyperlink to access the *Service-Instance-Name* Details page.
3. Click the **Passive Assurance** tab.
4. Expand the **L2 Circuit** accordion.

Table 236 on page 853 lists the fields in the Eline L2circuit NSM table.



NOTE: KPI-related data may not be displayed in the accordion tables and on the graphs when:

- There are device errors.
- The data displayed is derived from or is dependent on another KPI which is not being streamed by the devices.

To view health check results for the L2 circuit service connection status, click any hyperlinked entry in the Connection Status column of the Eline L2circuit NSM table. The corresponding accordion in the Eline L2circuit NSM Details for *service-instance-name* page displays the health check results.

Table 236: Fields in the Eline L2Circuit NSM Table

Field	Description
Device Name	Name of the PE device through which the L2 circuit service spans.
Neighbor	The IP address of the peer or neighbor PE device with which the PE device connects to create the L2 circuit in the MPLS network.
Connection ID	Unique ID for the L2 circuit connection between the two PE devices. The connection ID is in the <i>sub-interface-index (virtual-circuit-id)</i> format. For example, xe-0/0/0/0.0(vc 2).
Connection Status	The connection status of the L2 circuit: • Up—The L2 circuit connection is active or operational. • Down—The L2 circuit connection is not in active or operational state.

View Health Check Results for L2 Circuit Connection Status

The Eline L2circuit NSM Details for *service-instance-name* page displays the health check results for the connection status of the L2 circuit. You can view the PE device name, the IP address of the peer or neighbor PE device, and the L2 circuit connection ID at the top of the page. The accordion heading displays the parameter name, alerts generated for the parameter, and a condensed view of the performance graph displaying the health and functioning of the parameter.

To view the health check results in detail, expand the accordion. The accordion displays the performance graph for the L2 circuit connection status. The x-axis (horizontal) on the performance graph displays the time for which information is displayed. The graph displays a line with the data points related to events that have occurred. The line color indicates the connection status (Up or Down) of the L2 circuit. A blue line indicates that the status is Up and active. A gray line indicates that the status is Down and inactive. Click any data point on the line to view more information about the corresponding event in a pop-up. For more information about performance graphs, see "[Performance Graphs to View Passive Assurance Monitoring Data](#)" on page 871.

The peer or neighbor IP address and connection ID for the L2 circuit are displayed to the left of the graph in the *neighbor-IP-address: connection ID* format.

SEE ALSO

[Performance Graphs to View Passive Assurance Monitoring Data](#) | 871

[View Service Instance Details](#) | 836

View L3VPN Passive Assurance Monitoring Results

IN THIS SECTION

- [View BGP Health Check Results](#) | 856
- [View OSPF Health Check Results](#) | 858
- [View L3VPN Service Health Check Results](#) | 862

After provisioning an L3VPN service instance in the network, Routing Director automatically monitors the health of the service and its related parameters. You can view the health check results and the alerts

generated to indicate if any user action is required. The monitoring data is categorized into accordions and each accordion displays related information and alerts generated.

You can view the following accordions for L3VPN service instances:

- Physical Interfaces—View health check results for physical interfaces of PE devices through which the VPN service spans. Health check results include parameters such as:
 - Link flaps
 - Link state
 - Input and output bandwidth utilization
 - Input and output error counts
 - Interface speed

See ["Physical Interfaces Health Check Results" on page 864](#).

- Logical Interfaces—View health check results for operation status of the logical interfaces on the PE devices through which the VPN service spans. See ["Logical Interfaces Health Check Results" on page 869](#).
- BGP—View health check results for BGP neighbors. See ["View BGP Health Check Results" on page 856](#).
- OSPF—View health check results for OSPF neighbors. See ["View OSPF Health Check Results" on page 858](#).
- L3VPN—View health check results for L3VPN services. See ["View L3VPN Service Health Check Results" on page 862](#).



NOTE: KPI-related data may not be displayed in the accordion tables and on the graphs when:

- There are device errors
- The data displayed is derived from or is dependent on another KPI which is not being streamed by the devices

View BGP Health Check Results

IN THIS SECTION

- [View Health Check Results for BGP Parameters | 857](#)

The BGP accordion displays information about BGP peers and peering session states.

To access the BGP accordion:

1. Navigate to the **Orchestration > Instances > Service Instances** page.
2. Select a service instance and click the *service-instance-name* hyperlink to access the *Service-Instance-Name* Details page.
3. Click the Passive Assurance tab.
4. Click to expand the BGP accordion.

[Table 1 on page 856](#) lists the fields in the BGP Neighbors table.

To view health check results for the BGP neighbors, click any hyperlinked entry in the BGP Neighbors table to open the corresponding accordion in the BGP Detail page.

Table 237: Fields in the BGP Neighbors Table

Field	Description
Device Name	Name of the PE device on which the service is provisioned.
Peer Address	IP address of the peer or neighbor PE device in the MPLS network.

Table 237: Fields in the BGP Neighbors Table *(Continued)*

Field	Description
Peer State	Current peering session state of the BGP neighbor. The state of the neighbor: • Active—BGP is initiating a transport protocol connection in an attempt to connect to a peer. If the connection is successful, BGP sends an Open message. • Connect—BGP is waiting for the transport protocol connection to be completed. • Established—The BGP session has been established, and the peers are exchanging update messages. • Idle—BGP is waiting for a Start event and is the first stage of a connection. • OpenConfirm—BGP has acknowledged receipt of an open message from the peer and is waiting to receive a keepalive or notification message. • OpenSent—BGP has sent an open message and is waiting to receive an open message from the peer.
Flap Count	Number of BGP session flaps from the device.

View Health Check Results for BGP Parameters

The BGP Details page displays the health check results for the peering session state and session flap count of BGP neighbors, categorized into accordions. The accordion heading displays the parameter name, alerts generated for the parameter, and a condensed view of the performance graph displaying the health and functioning of the parameter.

To view the health check results in detail, click to expand the corresponding accordion. The accordion displays the performance graph for the peering session state and session flap count of BGP neighbors. For more information about performance graphs, see ["Performance Graphs to View Passive Assurance Monitoring Data" on page 871](#).

[Table 2 on page 858](#) lists the accordions on the BGP Detail page.

Table 238: Accordions on the BGP Detail Page

Accordion	Description
Peer State	View information and events related to peering session state of the BGP neighbors. The x-axis (horizontal) on the performance graph displays the time range for which information is displayed. The y-axis (vertical) displays the interface. The graph displays a line with the data points related to events that have occurred.
Flap Count	View information and events related to the number of times the BGP session flapped. The x-axis (horizontal) on the performance graph displays the time range for which information is displayed. The y-axis (vertical) displays the number of flaps. The graph displays a line with the data points related to events that have occurred.

View OSPF Health Check Results

IN THIS SECTION

- [View Health Check Results for OSPF Parameters | 861](#)

The OSPF accordion displays the information about OSPF neighbors related parameters.

To access the OSPF accordion:

1. Navigate to the **Orchestration > Instances > Service Instances** page.
2. Select a service instance and click the *service-instance-name* hyperlink to access the *Service-Instance-Name* Details page.
3. Click the Passive Assurance tab.
4. Click and expand the OSPF accordion.

[Table 3 on page 859](#) lists the fields in the OSPF Neighbor State table.

To view health check results for the OSPF neighbors, click any hyperlinked entry in the OSPF Neighbor State table to open the corresponding accordion in the OSPF Details page.

Table 239: Fields in the OSPF Neighbors State Table

Field	Description
Device Name	Name of the PE device on which the service is provisioned.
Instance Name	Name of the service instance configured on the PE device.
Neighbor Address	IP address of the interface that connects the PE device to a CE device.
Interface Name	Name of the interface of the PE device; for example, et-0/0/0.0.

Table 239: Fields in the OSPF Neighbors State Table *(Continued)*

Field	Description
Neighbor State	Displays the current peer session state of the OSPF neighbor. The state of the neighbor: • Attempt—Valid only for neighbors attached to nonbroadcast networks. It indicates that no recent information has been received from the neighbor, but that a more concerted effort must be made to contact the neighbor. • Down—Initial state of a neighbor conversation. It indicates that no recent information has been received from the neighbor. Hello packets might continue to be sent to neighbors in the Down state, although at a reduced frequency. • Exchange—Routing device is describing its entire LSDB by sending database description packets to the neighbor. Each packet has a sequence number and is explicitly acknowledged. • ExStart—First step in creating an adjacency between the two neighboring routing devices. The goal of this step is to determine which routing device is the primary, and to determine the initial sequence number (ISN). • Full—Neighboring routing devices are fully adjacent. These adjacencies appear in router link and network link advertisements. • Init—A hello packet has recently been sent by the neighbor. However, bidirectional communication has not yet been established with the neighbor. This state might occur, for example, because the routing device did not appear in the neighbor's hello packet. • Loading—Link-state request packets are sent to the neighbor to acquire more recent advertisements

Table 239: Fields in the OSPF Neighbors State Table *(Continued)*

Field	Description
	that have been discovered (but not yet received) in the Exchange state. • 2-Way—Communication between the two routing devices is bidirectional. This state has been ensured by the operation of the Hello Protocol. 2-Way is the most advanced state short of beginning adjacency establishment. The (backup) designated router is selected from the set of neighbors in state 2Way or greater.
Neighbor ID	IP address of the CE device.

View Health Check Results for OSPF Parameters

The OSPF Details page displays the health check results for the peering session state of OSPF neighbors. You can view PE device name, interface name, IP address of the CE device, instance name, and IP address of the interface that connects PE device to CE device at the top of the page. The accordion heading displays the parameter name, alerts generated for the parameter, and a condensed view of the performance graph displaying the health and functioning of the parameter.

To view the health check results of a parameter in detail, click and expand the accordion. The accordion displays the performance graph for the OSPF neighbor state. For more information about performance graphs, see ["Performance Graphs to View Passive Assurance Monitoring Data" on page 871](#).

[Table 4 on page 861](#) lists the accordion on the OSPF Detail page.

Table 240: Accordion on the OSPF Detail Page

Accordion	Description
Neighbor State	View information and events related to the peering session state of OSPF neighbors. The x-axis (horizontal) on the performance graph displays the time range for which information is displayed. The y-axis (vertical) displays the interface. The graph displays a line with the data points related to events that have occurred.

View L3VPN Service Health Check Results

IN THIS SECTION

- [View Health Check Results for L3VPN Parameters | 863](#)

The L3VPN accordion displays the information about L3VPN services.

To access the L3VPN accordion:

1. Navigate to the **Orchestration > Instances > Service Instances** page.
2. Select a service instance and click the *service-instance-name* hyperlink to access the *Service-Instance-Name* Details page.
3. Click the Passive Assurance tab.
4. Click and expand the L3VPN accordion.

[Table 5 on page 862](#) lists the fields that are displayed in the L3VPN table.

To view health check results for the L3VPN parameters, click any hyperlinked entry in the L3VPN table to open the corresponding accordion in the L3VPN Detail page.

Table 241: Fields in the L3VPN Table

Field	Description
VPN Network Access ID	Unique ID that you assigned to the site network access or connection between the PE device and the CE device.
VPN State	VPN connection status. You can view the VPN state for BGP and OSPF neighbors. For static and direct routes, the VPN state is not available.
Device Name	Name of the PE device used on which the L3VPN service is provisioned.

Table 241: Fields in the L3VPN Table (Continued)

Field	Description
Protocol Used	Name of the protocols used for the L3VPN service between the PE device and the CE device. Routing Director supports the OSPF and BGP routing protocols. For static and direct routes, the value is blank.
Neighbor Session	Peering session state of the neighbors. A session is formed between neighbors with the same AS in a network. A routing protocol such as BGP, OSPF runs between the established session to exchange routes.
Interface Name	Name of the interface on the PE device that connects to the CE device; for example, ge-0/0/0.
Instance IFL No	Logical interface instance of a PE device; for example, 100.
Interface Status	Status of the logical interface of a PE device.
Routing Instance Name	Name of the routing instance configured on the PE device.

View Health Check Results for L3VPN Parameters

The L3VPN Details page displays the health check results of the L3VPN parameters, categorized into accordions. You can view the name of the PE device on which service is provisioned. The accordion heading displays the parameter name, alerts generated for the parameter, and a condensed view of the performance graph displaying the health and functioning of the parameter.

To view the health check results of an interface parameter in detail, click to expand the corresponding accordion. The accordion displays the performance graph for the parameter. For more information about performance graphs, see ["Performance Graphs to View Passive Assurance Monitoring Data" on page 871](#).

Table 6 on page 864 lists the accordions on the L3VPN Detail page.

Table 242: Accordions on the L3VPN Detail Page

Accordion	Description
VPN State	View information and events related to the VPN state of the BGP or OSPF neighbors. For static and direct routes, the VPN State is not available. The x-axis (horizontal) on the performance graph displays the time range for which information is displayed. The y-axis (vertical) displays the VPN state. The graph displays a line with the data points related to events that have occurred.
Interface Status	View information and events related to the status of the interface from a PE device. The x-axis (horizontal) on the performance graph displays the time range for which information is displayed. The y-axis (vertical) displays the interface status. The graph displays a line with the data points related to events that have occurred.

Physical Interfaces Health Check Results

IN THIS SECTION

- [View Health Check Results for Physical Interface Parameters | 866](#)

The Physical Interfaces accordion displays the physical interface components on the PE devices that the VPN services span through.

To access the Physical Interfaces accordion:

1. Navigate to the **Orchestration > Instances > Service Instances** page.
2. Select a service instance and click the *service-instance-name* hyperlink to access the *Service-Instance-Name* Details page

3. Click the Passive Assurance tab.
4. Click and expand the Physical Interfaces accordion.

Table 243 on page 865 lists the fields in the Physical Interfaces table.



NOTE: KPI-related data may not be displayed in the accordion tables and on the graphs when:

- There are device errors.
- The data displayed is derived from or is dependent on another KPI which is not being streamed by the devices.

To view details of health check results for physical interface parameters, click any hyperlinked entry in the Physical Interfaces table to open the corresponding accordion on the Physical Interfaces Details for *service-instance-name* page.

Table 243: Fields in the Physical Interfaces Table

Field	Description
Device Name	Name of the PE device to which the physical interface belongs.
Interface Name	Name of the physical interface. For example, et-0/0/1.
Input Traffic Rate	The input traffic rate on the interface. Traffic rates will appear blank when devices do not send zero counters.
Output Traffic Rate	The output traffic rate on the interface. Traffic rates will appear blank when devices do not send zero counters.
Flaps	The number of link flaps on the interface.
Input Errors Count	The number of errors in packets received by the interface.
Output Errors Count	The number of errors in packets sent by the interface.

Table 243: Fields in the Physical Interfaces Table (Continued)

Field	Description
Admin State	The administrative state of the physical interface: • Up—The interface is active and can send and receive network traffic. • Down—The interface is inactive and does not send and receive network traffic. An interface is inactivated by the network administrator for reasons such as troubleshooting and maintenance.
Link State	The link state of the physical interface: • Up—The interface physical link is functional and network traffic can be sent or received by the interface. • Down—The interface physical link is not functional and the interface cannot transmit network traffic.
Speed	The interface speed or the maximum data transmission rate that the interface supports.

View Health Check Results for Physical Interface Parameters

The Physical Interfaces Details for *service-instance-name* page displays the health check results of the physical interface parameters, categorized into accordions. View the name of the physical interface for which the health check results are displayed and the name of the PE device to which the physical interface belongs at the top of the page. The accordion heading displays the parameter name, alerts generated for the parameter, and a condensed view of the performance graph displaying the health and functioning of the interface parameter.

To view the health check results of an interface parameter in detail, expand the corresponding accordion. The accordion displays the performance graph for the parameter. For more information about performance graphs, see ["Performance Graphs to View Passive Assurance Monitoring Data" on page 871](#).

[Table 244 on page 867](#) lists the accordions on the Physical Interfaces Details for *service-instance-name* page.

Table 244: Accordions on the Physical Interfaces Details Page

Accordion	Description
Input Utilization	View information and events related to input bandwidth utilization on the interface. The x-axis (horizontal) on the performance graph displays the time range for which information is displayed. The y-axis (vertical) displays the percentage of input bandwidth that is utilized. The graph displays a line with the data points related to events that have occurred. Click any data point to view more information about the corresponding event in a pop-up.
Output Utilization	View information and events related to output bandwidth utilization on the interface. The x-axis (horizontal) on the performance graph displays the time range for which information is displayed. The y-axis (vertical) displays the percentage of output bandwidth that is utilized. The graph displays a line with the data points related to events that have occurred. Click any data point to view more information about the corresponding event in a pop-up.
Flaps	View information and events related to interface link flaps. The x-axis (horizontal) on the performance graph displays the time for which information is displayed. The y-axis (vertical) displays the number of link flaps. The graph displays a line with the data points related to events that have occurred. Click any data point to view more information about the corresponding event in a pop-up.
Input Errors Count	View information and events related to errors in packets received by the interface. The x-axis (horizontal) on the performance graph displays the time for which information is displayed. The y-axis (vertical) displays the number of input errors. The graph displays a line with the data points related to events that have occurred. Click any data point to view more information about the corresponding event in a pop-up.
Output Errors Count	View information and events related to errors in packets sent by the interface. The x-axis (horizontal) on the performance graph displays the time for which information is displayed. The y-axis (vertical) displays the number of output errors. The graph displays a line with the data points related to events that have occurred. Click any data point to view more information about the corresponding event in a pop-up.

Table 244: Accordions on the Physical Interfaces Details Page *(Continued)*

Accordion	Description
Admin State	View information and events related to the administrative state of the interface. The performance graph displays the time range for which information is displayed on the x-axis and a line with data points related to events that have occurred. The line color indicates the state (Up or Down) of the interface. A blue line indicates that the status is Up or in active state. A gray line indicates that the status is Down or inactive. Click any data point on the line to view more information about the corresponding event in a pop-up. The interface name is displayed to the left of the graph.
Link State	View information and events related to the link state of the interface. The performance graph displays the time range for which information is displayed on the x-axis and a line with data points related to events that have occurred. The line color indicates the link state (Up or Down) of the interface. A blue line indicates that the link state is Up or active. A gray line indicates that the link state is Down or inactive. Click any data point on the line to view more information about the corresponding event in a pop-up. The interface name is displayed to the left of the graph.
Speed	View information and events related to the speed of the interface. The x-axis (horizontal) on the performance graph displays the time for which information is displayed. The y-axis (vertical) displays the speed (in Mbps) of the interface. The graph displays a line with the data points related to events that have occurred. Click any data point to view more information about the corresponding event in a pop-up.

SEE ALSO

[Performance Graphs to View Passive Assurance Monitoring Data | 871](#)

[View Service Instance Details | 836](#)

Logical Interfaces Health Check Results

IN THIS SECTION

- [View Health Check Results for Logical Interface Operation Status | 870](#)

The Logical Interfaces accordion displays the components that are related to the logical interfaces on the PE devices that the VPN services span through.

To access the Logical Interfaces accordion:

1. Navigate to the **Orchestration > Instances > Service Instances** page.
2. Select a service instance and click the *service-instance-name* hyperlink to access the *Service-Instance-Name* Details page
3. Click the Passive Assurance tab.
4. Click and expand the Logical Interfaces accordion.

[Table 245 on page 869](#) lists the fields in the Logical Interfaces table.

To view details of health check results for logical interface operational status, click any hyperlinked entry in the Logical Interfaces table to open the corresponding accordion on the Logical Interfaces Details for *service-instance-name* page.

Table 245: Fields in the Logical Interfaces Table

Field	Description
Device Name	Name of the PE device associated with the logical interface. The logical or virtual interface is created on a physical interface of the device.
Interface Name	Name of the physical interface on which the logical interface is created. For example, et-0/0/1.
Interface IFL No	Logical Interface number of a PE device; for example 100.

Table 245: Fields in the Logical Interfaces Table (*Continued*)

Field	Description
Input Traffic Rate	The input traffic rate on the interface. Traffic rates will appear blank when devices do not send zero counters.
Output Traffic Rate	The output traffic rate on the interface. Traffic rates will appear blank when devices do not send zero counters.
Operation Status	The operational status of the logical interface: • Up—The interface is active or operational and can send and receive network traffic. • Down—The interface is not in active or operational state.

View Health Check Results for Logical Interface Operation Status

The Logical Interface Details for *service-instance-name* page displays the health check results for the operational status of the logical interface. You can view the name of the PE device, the name of the physical interface on which the logical interface is created, and the logical or subinterface index at the top of the page. The accordion heading displays the parameter name, alerts generated for the parameter, and a condensed view of the performance graph displaying the health and functioning of the parameter.

To view the health check results in detail, expand the accordion. The accordion displays the performance graph for the logical interface operation status. The x-axis on the performance graph displays the time for which information is displayed. The graph displays a line with clickable data points related to events that have occurred. The line color indicates the operation status (Up or Down) of the logical interface. A blue line indicates that the status is Up and active. A gray line indicates that the status is Down and inactive. Click any data point on the line to view more information about the corresponding event in a pop-up. For more information about performance graphs, see ["Performance Graphs to View Passive Assurance Monitoring Data" on page 871](#).

The physical and subinterface or logical interface names are displayed to the left of the graph in the *physical-interface-name : subinterface index* format.

SEE ALSO

[Performance Graphs to View Passive Assurance Monitoring Data | 871](#)

[View Service Instance Details | 836](#)

Performance Graphs to View Passive Assurance Monitoring Data

The Passive Assurance tab in the *Service-Instance-Name* Details page allows you to view results of health checks that Routing Director performs on a service and its related parameters after the service is provisioned in the network. To access the Passive Assurance tab, navigate to the **Orchestration > Instances > Service Instances** page, and select and click the service instance name to open the *Service-Instance-Name* Details page. Click the Passive Assurance tab to view the service components and related parameters, categorized into accordions. Each accordion under the Passive Assurance tab provides information about the related parameters and the relevant events (alerts) generated. Click any hyperlinked table entry to open the corresponding accordion on the *Service-Component-Name* Details for *service-instance-name* page.

The *Service-Component-Name* Details for *service-instance-name* page lists the parameters (categorized into accordions) for which Routing Director performs health checks. The performance graph for an accordion displays information about the performance of the parameter in a specified time duration.

Performance graphs for measurable parameters such as bandwidth, flaps, error counts, and so on display X (horizontal) and Y (vertical) axes. The x-axis shows the time duration for which the performance is displayed. The y-axis displays the parameter name and number or measurement unit. The graphs display a line with the data points related to events that have occurred. Click any data point to view more information about the corresponding event in a pop-up. The cause for the alert, if any, is also displayed.

Performance graphs for the operational or connection status of a parameter display only the x-axis that shows the time duration for which the performance is displayed. The graphs display a blue or gray line with data points to indicate whether the status is Up or Down. A blue line indicates that the parameter status is Up or in active state. A gray line indicates that the status is Down or not operational. Click any data point on the line to view more information about the corresponding event in a pop-up. The cause for the alert, if any, is also displayed.

You can view the performance of a parameter for different time durations or set a custom duration. By default, information for the past 30 minutes is displayed. To change this period, click one of following options provided in the time duration bar on top of the *Service-Component-Name* Details for *service-instance-name* page:

- **Week**
- **Day**
- **3 hrs**

- 1 hr
- 30 mins
- Custom

Click the Custom button to set a custom time duration for which you want to view the performance graph. The custom time setting bar is enabled and displays the start and end time for which you are viewing the performance graphs in the MM/DD/YYYY HH:MM AM/PM format. Click the start and end time to change the time duration. The calendar is displayed and you can select the new start and end time duration to view the performance graphs.



NOTE: The date you set must be the present date or any date from the past one week.

You can refresh the performance graphs to view the latest data by clicking the Refresh icon at the top of the *Service-Component-Name* Details for *service-instance-name* page. The performance graphs also auto-refresh at an interval depending on the time range for which the information is displayed. [Table 246 on page 872](#) lists the auto-refresh frequencies for different time ranges. The next refresh time is displayed near the Refresh icon at the top of the page.

You can also view the performance graphs for service component parameters in a new tab. Click the pop-out icon at the top-right corner of the accordion to view the performance graph in a new tab.

Table 246: Auto-Refresh Frequency for Performance Graphs

Time Range	Auto-Refresh Frequency
Weekly	Every 16.8 hours
Daily	Every 58 minutes
Every three hours	Every 8 minutes
Hourly	Every 3 minutes
Thirty minutes	Every 2 minutes
Custom	No auto-refresh

RELATED DOCUMENTATION

[Physical Interfaces Health Check Results | 864](#)

[Logical Interfaces Health Check Results | 869](#)

[View L2 Circuit Passive Assurance Monitoring Data | 852](#)

Monitor Service Order Execution Workflows

IN THIS CHAPTER

- [About the Service Orders Page | 874](#)
- [About the Workflows Page | 880](#)
- [About the Workflow Runs Page | 882](#)
- [View Workflow Run Details | 885](#)

About the Service Orders Page

IN THIS SECTION

- [Tasks You Can Perform | 875](#)
- [Field Descriptions | 876](#)

The Service Orders page is an inventory of service orders that Routing Director generates for provisioning or deleting network services. A service order specifies the type of operation that you want to implement for a service instance. Routing Director generates service orders when you perform the following operations:

- Create a service instance—A create service order is generated when you create a service instance.
- Modify a service instance—A modify service order is generated when you modify an existing service instance.
- Delete a service instance—A delete service order is generated when you delete or deprovision a service instance.

You can create or modify a service instance, save and provision the instance. A corresponding service order is generated to provision the new or modified service in the network. Deleting a service by

selecting the service instance and clicking **More > Deprovision** on the Service Instances page generates a delete service order.

Each type of service order has a corresponding workflow that Routing Director activates after you provision a service instance. For example, when you provision a newly added service instance, Routing Director activates the create workflow to provision the service. Similarly, the modify and delete workflows are activated when you provision a modified service instance and delete a service instance respectively. Routing Director processes only a single service order for a service instance at any given time. If you provision more than one service order for a service instance at a given time, the additional service orders are queued.

Routing Director generates service orders for:

- Onboarding and offboarding devices
- Creating, modifying, or deleting network services
- Creating, modifying, or deleting resource pools for auto allocation when provisioning services

The Service Orders page lists all service orders that Routing Director generates. You can view details such as the service order name, operation type, execution state, and so on.

To access the Service Orders page, click **Orchestration > Service Orders**.

Tasks You Can Perform

- View a list of all service orders that you initiated. See [Table 247 on page 876](#).
- View details of each service order—To view details of each service order, click the **Details** icon that appears next to the service order name when you hover on the service order. The *service-order-name* pane appears on the right side of the page. The pane lists details about the service order and includes information such as the order properties and status.

You can also view the service order in the JSON format by clicking **More > JSON View**. See [Table 248 on page 878](#).

Additionally, click **Download** to download the service order in the JSON format to your local system.



NOTE: The Order Status section of the *service-order-name* pane displays information only if the service order is provisioned. If a service order is uploaded but not provisioned, the Order Status does not display any information.

- Filter the data displayed in the table—Click the filter icon (funnel) and select whether you want to show or hide advanced filters. You can then add or remove filter criteria, save criteria as a filter, apply or clear filters, and so on. The filtered results are displayed on the same page.

- Search by using keywords—Click the search icon (magnifying glass), enter the search term in the text box, and press Enter. The search results are displayed on the same page.
- Sort, resize, or re-arrange columns in a table (grid).
- Show or hide columns in the table or reset page preferences, using the vertical ellipsis menu.
- Set the number of rows to be displayed in the table—Click the **Display** drop-down at the bottom-right corner of the Service Orders table to set the number of rows (15, 30, 50, or 100) to be displayed in the table. The default number of rows displayed is 15.

Use the arrow heads to traverse through the pages or enter a page number to view the entries listed on that page.

- Set auto refresh frequency—Click the Refresh icon at the bottom-left corner of the Service Orders table to view the most recent updates.

You can also auto refresh the table to view updated details. To set auto refresh frequency:

1. Click the horizontal ellipsis near the Refresh icon.
2. Select the refresh frequency from the available options (30 seconds, 1 minutes, 3 minutes, and 5 minutes).

The next refresh time is displayed near the Refresh icon.

To turn off auto refresh, click the horizontal ellipsis and click **Off**.

Field Descriptions

Table 247 on page 876 lists the fields on the Service Orders page.

Table 247: Fields on the Service Orders Page

Field	Description
Name	Name of the service order.
Order ID	Unique alphanumeric ID assigned to a service order.
Customer	Name of the customer for whom the service order is executed.

Table 247: Fields on the Service Orders Page (*Continued*)

Field	Description
Operation	Type of operation for which the service order is executed. Service orders are executed when you perform the following operations for a service instance: • Create • Modify • Delete
Service Design	Name of the service design associated with the service order.
Design Version	Service design version associated with the service order.
State	State of execution of the service order. A service order can have the following execution states: • Success—The service order is successfully executed. • Processing—The service order is in the process of being executed. • Failed—The service order has failed to execute. • Success with warnings—The service order is successfully executed but certain tasks in the workflow have failed. You can view the logs of tasks that fail and take further action if required. For information about viewing task logs, see "View Workflow Run Details" on page 885.
Status	The status of all service orders that you upload and provision. For example, for a service order that is uploaded but not yet provisioned, the status displays as Uploaded. For a device onboarding service order that is executed, the status displays as Applied configuration to <i>device-name</i>.

Table 247: Fields on the Service Orders Page (Continued)

Field	Description
Workflow Run ID	Unique ID assigned to a workflow run associated with a service order. Click the workflow run ID hyperlink associated with a service order to view details about the workflow run on the Workflow Runs List page. See "View Workflow Run Details" on page 885.
Upload Time	The date and time when the service order was generated.
Edited By	Name of user who provisioned the service instance that generated the service order.

Table 248: Fields on the Service Order Details Pane

Field	Description
Properties	
Name	Name of the service order.
Status	Status of the service order.
Order Id	Unique alphanumeric ID assigned to the service order.
Last Modified	Date and time when the service instance related to the service order was last modified.
Last Modified By	Name of the user who last modified the service instance associated with the service order.
Order Status	

Table 248: Fields on the Service Order Details Pane (*Continued*)

Field	Description
Status	State of execution of the service order: • Success • Processing • Failed • Success with warnings
Components	Lists the components of a service order and the tasks executed as part of the service order workflow, for each component: • Placement—Click placement to view the placement-specific workflow and task execution status. Placement refers to the allocation of resources to implement a service. These resources includes site locations, ports and interfaces, VLANs, CEs, and so on. • Device—Click device to view the device-specific workflow and task execution status. • Insights—Click insights to view the Insights (passive assurance) configuration workflow and task execution status. • Active Assurance—Click active assurance to view the Active Assurance configuration workflow and task execution status. To view details about the workflow run tasks associated with a component, click the View Workflow Run hyperlink below the component heading. The Workflow Runs List opens and you can view details about the workflow runs and task logs. See "View Workflow Run Details" on page 885.
Workflow	Click Workflow to view the complete workflow and task execution status of the service order. To view details about the workflow run tasks, click the View Workflow Run hyperlink. The Workflow Runs List opens and you can view details about the workflow runs and task logs. See "View Workflow Run Details" on page 885.

RELATED DOCUMENTATION

[About the Workflow Runs Page | 882](#)

[View Workflow Run Details | 885](#)

[About the Service Instances Page | 829](#)

About the Workflows Page

IN THIS SECTION

● [Tasks You Can Perform | 880](#)

● [Field Descriptions | 881](#)

Routing Director uses automated workflows or series of tasks to provision a service in the network. A workflow is generated for each service order to create, modify, or delete a service instance in the network. A workflow run refers to an instance of executing the tasks in a workflow. You can monitor the status of workflow execution from the Workflows page. Use this page to view details such as workflow execution states, workflow run status, and the details of the last workflow run.



NOTE: The Workflows page lists data for up to 1000 workflow runs that were executed in the last 90 days.

To access the **Workflows** page, click **Orchestration > Workflows** on the navigation menu.

Tasks You Can Perform

You can perform the following tasks on the Workflows page:

- View workflow details such as workflow IDs, states of workflow execution, time and status of the last workflow run, and so on. See [Table 249 on page 881](#).

Click the **All Workflow Runs** hyperlink at the top of the page to open the Workflow Runs page and view the list of all workflow runs. By default, the workflow runs are sorted by the run execution time, with the latest runs being listed on top of the list. You can change the view and filter runs by the execution states. Click the drop-down under the State label and select the execution state by which you want to filter the workflow runs.

Clicking a ***workflow-id*** opens the workflow runs list page. You can view all runs executed for the workflow and the details of each run from the workflow runs list page. See ["View Workflow Run Details" on page 885](#).

- Set auto-refresh frequency—Click the Refresh icon at the bottom-left corner of the Workflows table to view the most recent updates.

You can also auto refresh the table to view updated details. To set auto refresh frequency:

1. Click the horizontal ellipsis near the Refresh icon.
2. Select the refresh frequency from the available options (30 seconds, 1 minutes, 3 minutes, and 5 minutes).

The next refresh time is displayed near the Refresh icon.

To turn off auto refresh, click the horizontal ellipsis and click **Off**.

- Show or hide columns in the table or reset page preferences, using the vertical ellipsis menu.
- Sort, resize, or re-arrange columns in a table (grid).

Field Descriptions

[Table 249 on page 881](#) lists the fields on the Workflows page.

Table 249: Fields on the Workflows Page

Field	Description
Workflow ID	Unique ID generated for a workflow.
Description	Description about the workflow, for example, create L3VPN service.
Owners	Name of user who owns the workflow.
Queued	Number of workflow runs that are queued for execution. Click the number to view the run and task details.
Success	Number of workflow runs that were successfully executed. Click the number to view each run and the associated task details.

Table 249: Fields on the Workflows Page *(Continued)*

Field	Description
Running	Number of workflow runs that are currently being executed. Click the number to view each workflow run and the associated task details.
Failed	Number of workflow runs that have failed to execute. Click the number to view each workflow run and the associated task details.
Last Run	The date and time of the last workflow run.
Last Run Status	The status of the last workflow run: • Success—The last workflow run was successfully executed. • Failed—The last workflow run failed to execute.
Schedule	The date and time when the service order for a service instance is scheduled.
Next Run	The date and time when the next workflow run is scheduled.

RELATED DOCUMENTATION

[About the Workflow Runs Page | 882](#)

[View Workflow Run Details | 885](#)

About the Workflow Runs Page

IN THIS SECTION

- [Tasks You Can Perform | 883](#)
- [Field Descriptions | 884](#)

Use the Workflow Runs page to view details of workflow runs for service orders generated to create, modify, or delete a service instance. A workflow run ID is generated and assigned to each workflow run. The Workflow Runs page lists the runs that have been initiated and the execution state of each run.



NOTE: The Workflow Runs page lists data for up to 1000 workflow runs that were executed in the last 90 days.

To access the Workflow Runs page:

1. Click **Orchestration > Workflows** on the navigation menu.
2. Click a hyperlinked number in any of the workflow execution state columns (Queued, Success, Running, Failed).

You are directed to the Workflow Runs page, where the state is used as a filter to display the runs in that state. To remove filters and view the list of all workflow runs, click **X** on the right side of the filtering bar.

Tasks You Can Perform

You can perform the following tasks on the Workflow Runs page:

- View a summary of the workflow runs by the run ID, workflow ID, state of the run, and the start date and time and the end date and time of the run. See [Table 250 on page 884](#).
- View the workflow run details and logs for the associated tasks.

To view the details of the workflow runs, click the *workflow ID* or *run ID* hyperlink.

Clicking the *run ID* opens the Workflow Runs List page. The selected run is highlighted and details about the run are displayed on the right side of the page.

- Filter the data displayed in the table—Click the filter icon (funnel) and select whether you want to show or hide advanced filters. You can then add or remove filter criteria, save criteria as a filter, apply or clear filters, and so on. The filtered results are displayed on the same page.
- Show or hide columns in the table or reset page preferences, using the vertical ellipsis menu.
- Sort, resize, or re-arrange columns in a table (grid).
- Set auto refresh frequency—Click the Refresh icon at the bottom-left corner of the Workflow Runs table to view the most recent updates.

You can also auto refresh the table to view updated details. To set auto refresh frequency:

1. Click the horizontal ellipsis near the Refresh icon.

2. Select the refresh frequency from the available options (30 seconds, 1 minutes, 3 minutes, and 5 minutes).

The next refresh time is displayed near the Refresh icon.

To turn off auto refresh, click the horizontal ellipsis and click **Off**.

Field Descriptions

Table 250 on page 884 lists the fields on the Workflow Runs page.

Table 250: Fields on the Workflow Runs Page

Field	Description
State	State of the workflow run: • Success—The workflow run is successfully executed. • Failed—The workflow run failed to execute.
Workflow ID	Unique ID of the workflow for which the run is being executed.
Time	The date and time when the workflow run was executed.
Run ID	Unique ID assigned to the workflow run.
Service Instance	Name of the service instance for which the service order workflow run is executed. To view details about the service instance, click the <i>service instance name</i> hyperlink. You are directed to the Service Instances page where the instance name is used as a filter to display details about the service instance.
Service Order Reference	The service order ID for which the workflow run is executed. The order ID is a unique alphanumeric number assigned to a service order. To view details about the service order, click the <i>order ID</i> hyperlink. You are directed to the Service Orders page where the order ID is used as a filter to view details about the service order.

Table 250: Fields on the Workflow Runs Page *(Continued)*

Field	Description
Start Date	The date and time when a workflow run started to execute. The start date and time is specified in the workflow.
End Date	The date and time when a workflow run stopped execution.
Note	Additional information about the run.

RELATED DOCUMENTATION

[About the Workflows Page | 880](#)

[View Workflow Run Details | 885](#)

View Workflow Run Details

IN THIS SECTION

- [View Task Logs | 887](#)

Routing Director allows you to view details of workflow runs and troubleshoot errors by viewing detailed task logs. To access the Workflow Runs List page, click **Orchestration** > **Workflows** on the navigation menu and use one of the following options on the Workflows page:

- Click the *workflow ID* for which you want to view the run details.
- Click a hyperlinked number in any of the workflow execution state columns (Queued, Success, Running, Failed). You are directed to the Workflow Runs page, where the state is used as a filter to display the runs in that state. On the Workflow Runs page, click the *workflow ID* or *run ID* for which you want to view details. You are directed to the Workflow Runs List page.



NOTE: The Workflow Runs List page lists data for up to 1000 workflow runs that were executed in the last 90 days.

You can view the following details about the selected run:

- The **Task** tab displays the list of tasks and detailed logs for each task. See [Table 252 on page 887](#).
- The **Graph** tab displays a graphical representation of the workflow tasks that are executed as part of the run. Click **Vertical** for a vertical view of the workflow chart. Click **Horizontal** for a horizontal view of the flowchart.
- The **Details** tab displays general details about the run such as the execution status, run ID and duration, workflow ID, and so on. See [Table 251 on page 886](#).

You can refresh the Workflow Runs List page by clicking the refresh icon. You can also set auto refresh frequency to refresh the page by clicking the horizontal ellipsis near the refresh icon and selecting the refresh frequency (30 seconds, 1 minute, 3 minutes, or 5 minutes).

Table 251: Fields on the Details Tab

Field	Description
Status	The status of the workflow run: • Success—The workflow run was successfully executed. • Failed—The workflow run failed to execute.
Run ID	Unique ID assigned to the workflow run.
Run Duration	Duration of time taken to execute the workflow run.
Logical Date	The start of the data interval for a workflow run. The data interval refers to the time range in which the workflow run is scheduled. For example, for a workflow run that is scheduled hourly, the data interval is an hour.
Start Date	The date and time when a workflow run started to execute. The start date and time is specified in the workflow.
End Date	The date and time when a workflow run stopped execution.

Table 251: Fields on the Details Tab *(Continued)*

Field	Description
Workflow ID	Unique ID generated for a workflow.
Configuration	Details associated with the workflow run, including the service instance ID, service order name and operation type, customer and organization IDs, instance path, and lock ID.

Table 252: Fields on the Task Tab

Field	Description
Task	Name of the task. You can view detailed logs of a task by clicking the details icon near the task name. See "View Task Logs" on page 887 .
State	Task execution status: • Success—The task was successfully executed. • Skipped—The task was skipped during workflow run execution. • Failed—The task failed to execute. • Upstream_failed—The current task failed to execute as its upstream or preceding task on which it has a dependency failed to execute.
Started	Date and time when the task was started.
Duration	Duration of time taken to execute the task.

View Task Logs

You can view and copy detailed logs of tasks for troubleshooting and debugging. Use one of the following ways to view task logs from the **Task** tab:

- Click the **Details** icon that appears next to the task name when you hover over a task.
- Select the task and click **More > Detail**.

The task details pane appears and lists the logs for the selected task. You can perform the following tasks on the details pane:

- Filter logs by log levels—Click the **Log by Attempt** drop-down to filter the logs by the following log levels:
 - Debug
 - Info
 - Warning
 - Error
 - Critical
- View logs in the word wrap format by selecting the **Word Wrap** check box.
- View log timestamps in the ISO format by selecting the **ISO Format** check box.
- Copy and save the logs on your system for troubleshooting or later use. Click **Copy** and save the logs.

RELATED DOCUMENTATION

[About the Workflows Page | 880](#)

[About the Workflow Runs Page | 882](#)

7

PART

Active Assurance

- Introduction | **890**
 - Test Agents | **897**
 - Tests and Monitors | **944**
 - Test Templates | **1288**
-

Introduction

IN THIS CHAPTER

- [Active Assurance Overview | 890](#)
- [Active Assurance Workflow | 891](#)
- [Active Assurance Terminologies | 892](#)

Active Assurance Overview

IN THIS SECTION

- [Benefits | 891](#)

Active Assurance uses active, synthetic traffic to continuously monitor your network, which in turn helps you to understand your network quality, and identify, troubleshoot, and resolve issues before customers notice them.

Traditional service assurance only collects telemetry data from devices in the infrastructure and are not designed to determine whether a service is properly working from the end user's perspective. Furthermore, traditional assurance solutions do not assess the customer experience over the lifetime of the service. With Active Assurance, you can automate testing processes and gain continuous end-to-end service quality insights to proactively enhance the customer experience.

The following are the key components of Active Assurance:

- **Control Center—Routing Director** acts a control center, which manages Test Agents that runs on devices in your network. After you install Routing Director in your premises, you can log in to the Routing Director GUI and manage Test Agents in your network. For more information on installation instructions, see *Routing Director Installation Guide*.

- **Test Agent**—Software installed on network devices that generate and receive traffic from one or more Test Agents and receive control information from Routing Director.
- **Plug-ins**—Software feature that executes the traffic generation when a Test or a Monitor is created. The Test Agent downloads the plug-in that are executable from Routing Director.

You can access the Routing Director GUI to view the Test Agents running across your network, create and run on-demand Tests and Monitors, and view real-time and aggregated results of these Tests and Monitors to gain insight into your network.

Benefits

- Provides real-time actionable insights into your network's health, performance, and quality. Active Assurance validates the data plane network by generating synthetic traffic from different points in the network. The packet ingestion in the network reflects the network's behavior from multiple physical points using different protocols. This approach validates the network's configuration behavior from the end-user or from the application perspective.
- Presents a consolidated user interface for managing all Test Agents in your network and for viewing real-time aggregated views of Monitors and Tests.

RELATED DOCUMENTATION

[Active Assurance Workflow | 891](#)

[Test Agents Overview | 897](#)

[Tests and Monitors Overview | 944](#)

Active Assurance Workflow

Active Assurance enables you to set up and run on-demand tests using Test Agents. Test Agents are measurement points, which are deployed in your network. These Test Agents are capable of generating, receiving, and analyzing network traffic and therefore enable you to continuously view and monitor both real-time and aggregated result metrics

The following is the workflow for monitoring your network using Active Assurance:

1. Install Routing Director. For information on installing Routing Director, see *Routing Director Installation Guide*.
2. Login in to the Routing Director GUI. See ["Administration Workflow" on page 55](#).

- 3. Onboard devices. For information on onboarding the devices, see ["Device Onboarding Workflow" on page 173](#).
- 4. Ensure that you install Test Agent on your devices so that you can manage these Test Agents using Routing Director. For information on Test Agents, see ["Test Agents Overview" on page 897](#) and ["Install Test Agent Application " on page 929](#).
- 5. Create Tests and Monitors using Measurement Designer. See ["Create a Test" on page 958](#) and ["Create a Monitor" on page 1110](#).
- 6. To continuously monitor the key performance indicators (KPIs), run Tests and Monitors that you have created based on your requirement.
- 7. View the results of streams and alerts generated.

See ["View Stream Details" on page 1258](#).
- 8. Analyze the metrics, identify the root cause of an issue (if any), and take necessary actions.

RELATED DOCUMENTATION

| [Active Assurance Overview | 890](#)

Active Assurance Terminologies

To understand Active Assurance, you should be familiar with the terms defined in [Table 253 on page 892](#).

Table 253: Terminologies Used in Active Assurance

Term	Definition
Test Agent	Test Agent is a software that is installed on your network device, and it acts as a measurement point in your network. These Test Agents are deployed at strategic locations in your network to evaluate the quality of your network by collecting metric data for preconfigured key performance indicators (KPIs). Test Agents generate, receive, and analyze network traffic, and therefore enable you to continuously view and monitor both real-time and aggregated result metrics.

Table 253: Terminologies Used in Active Assurance *(Continued)*

Term	Definition
Test Agent Application	Test Agent Application is a software component that can be installed within your network and delivered through a Docker container to run Tests and Monitors. Test Agent Application can be deployed in two ways: • On network devices—Test Agent Application is supported in Junos OS Evolved ACX routers. You can configure a Test Agent Application during onboarding a device from (Inventory > Device and Interface Profiles). You can also manage the Test Agent Applications by using APIs. • As a standalone Test Agent Application—The Test Agent Application is delivered as a Docker container image by following standard Docker deployment procedures. You can deploy the application in cloud-based or virtualized infrastructures, without the need for direct installation on physical devices.
Test Agent Appliance	A Test Agent Appliance is a full-fledged Test Agent with built-in operating system, which allows full control over network configuration and supports advanced functionalities. The Test Agent Appliance is based on Debian Linux operating system. Test Agent Appliance can be deployed in two ways: • As a Dedicated Test Agent—You can download a Test Agent Appliance Software image and install on a custom x86 hardware. • As a Test Agent Virtualized Network Function (TA VNF)—You can upload a Test Agent Software image to a virtualization platform and run on a virtual machine (VM) running on a hypervisor.

Table 253: Terminologies Used in Active Assurance *(Continued)*

Term	Definition
Plug-in	A plug-in is an extension to the Test Agent. The Test Agents download the plug-in from Routing Director, when required. The plug-in sends the actual traffic and performs the work that is specified in the Tasks and Monitors. Example of plug-in are HTTP, DNS, TCP, UDP, and so on.
Test	A Test is a set of verifications that is performed by one or more Test Agents for a finite amount of time. A Test contains of one or more Steps and Tasks. A Test can consist of several Steps that are executed sequentially. Each Step can consist of one or more Tasks that run concurrently. A Test delivers a binary result—Pass or Fail.
Monitor	A Monitor is a set of verifications that is performed by one or more Test Agents for an infinite amount of time. A Monitor can contain multiple Tasks, which run in parallel and can continuously monitor the KPIs that are defined in the Monitor. A Monitor delivers a time-framed result as it runs continuously.
Task	A Task is a part of Tests and Monitors. When you run a Test or a Monitor, a single task sends a set of Measurements to the Test Agent. The Task also has metrics Evaluation Criteria, which is passed on to the underlying Measurements.
Step	A Step runs a set of parallel Tasks. A Test can have one or more Steps that are executed sequentially, while a monitor has a single Step.

Table 253: Terminologies Used in Active Assurance *(Continued)*

Term	Definition
Evaluation Criteria	Configure customized evaluation criteria for metrics by using threshold expressions. On violation of any configured expression, Routing Director generates events.
Measurement	When you run a Test or a Monitor, Routing Director instructs the Test Agents to send or receive traffic using the selected protocols through measurements. Each measurement produces one or more Streams of metrics.
Stream	Stream represents the stream of data between a Test Agent and an endpoint in a single direction. Depending on the protocol used, an endpoint could be a webserver (HTTP), or another peer Test Agent (TCP). The Tests or Monitors evaluate these Streams and the summary of the Test or Monitor results is displayed on the Test Details (Observability > Active Assurance > Tests > <i>Test- Name</i>) or Monitor Details (Observability > Active Assurance > Monitors > <i>Monitor-Name</i>) pages. By looking at these Stream graphs, you can identify what caused the violation.
Errored second	An Errored Second is an aggregation value that is stored as a time-series metric as part of a Stream. It contains information about the number of seconds that the error (threshold violations) has occurred.

Table 253: Terminologies Used in Active Assurance (*Continued*)

Term	Definition
Tags	You can configure Tags for a Test or a Monitor in the <i>key:value</i> format. A tag is a key-value pair in which the key signifies a category for which you configure a value. The value is an identifier for the category. Examples of key-value pairs are device and device name (edgedevice:acx7000), site and site name (site:bangalore).

Test Agents

IN THIS CHAPTER

- [Test Agents Overview | 897](#)
- [About the Test Agents Page | 900](#)
- [Add a Test Agent | 908](#)
- [About the *Test Agent-Name* Page | 911](#)
- [Install Test Agent Application | 929](#)
- [Install Test Agent Appliance | 934](#)

Test Agents Overview

IN THIS SECTION

- [How Test Agents Measure Metrics | 898](#)
- [Benefits of Using Test Agents | 900](#)

Active Assurance uses Test Agents, which is a software that is installed on your network device and it acts as a measurement point in your network. These Test Agents are deployed at strategic locations in your network to evaluate the quality of your network by collecting metric data for pre-configured key performance indicators (KPIs). Test Agents generate, receive, and analyze network traffic and therefore enable you to continuously view and monitor both real-time and aggregated result metrics.

All Test Agents are controlled and updated remotely through Routing Director GUI.

Test Agents can be installed on Junos OS Evolved routers, x86 hardware, or on virtual machines through docker container.

For ACX7000 series routers, which uses Junos OS Evolved, the Test Agents can either be run as a docker container, or, the Test Agent is automatically installed during the onboarding process. For

ACX7000 series routers, you can also use the Real-Time Performance Monitoring (RPM) functionality to generate traffic and collect the metrics that help to analyze the network's behavior.

For MX and PTX series routers, you can use the Real-Time Performance Monitoring (RPM) functionality to generate traffic and collect the metrics that help to analyze the network's behavior. RPM enables you to configure active probes to track and monitor traffic. RPM functionality can be used with RPM plug-ins. When RPM is configured on a device, the device calculates network performance based on packet response time, jitter, and packet loss.

To gather RPM statistics, the device sends out probes to a specified probe target that is identified by an IP address. On receiving a probe, the target generates a response that is sent back to the device. See [Monitoring Traffic Using Real-Time Performance Monitoring \(RPM\)](#).



NOTE:

- When running concurrent RPM probes on a router, you must specify the probe limit between 1 through 500 to avoid the probe limit error. See [Limiting the Number of Concurrent RPM Probes](#).
- When you use RPM plug-ins, you must set the system-level timezone in the routers to manage daylight saving time (DST) effectively. In the absence of a proper timezone setting, RPM plug-ins rely on a fallback mechanism that checks the time difference from UTC every hour to adjust the offset.

In addition, you must synchronize routers with a Network Time Protocol (NTP) server for accurate timekeeping.

For any other x86 hardware or virtual machines (VMs), you need to manually install Test Agents. For information on installing Test Agents, see ["Install Test Agent Application " on page 929](#).


You can also use APIs to create, install, and uninstall test agents in Routing Director. See *API Reference Guide*.

The Test Agents in your network are listed on the Test Agents page (**Inventory > Active Assurance**). For more information on the Test Agents page, see ["About the Test Agents Page" on page 900](#).

How Test Agents Measure Metrics

[Figure 47 on page 899](#) illustrates the relationship between a Test Agent plug-in, measurements, tasks, and streams.

Figure 47: How Test Agents Create Streams



Routing Director leverages the capability of monitoring and testing the network's data plane by generating synthetic traffic using Test Agents or devices deployed in the network. Test Agents are configured to run a set of Measurements as defined by the Tests and Monitors.

Routing Director contains a set of features to produce Measurements in the network's data plane in the form of plug-ins (For example, Ping, HTTP, and so on). When you configure a Measurement using the Routing Director GUI or REST APIs, Routing Director configures the Test Agent to download the plug-in and pushes the configuration to Test Agent.

While configuring a Measurement, you can also set evaluation criteria. The evaluation criteria contain thresholds for the individual metrics. Routing Director uses the thresholds, which you have defined, to determine the network's behavior. If the threshold is violated, then Routing Director generates events.

A Measurement produces one or more Streams. The Stream contains the measured metrics (also known as KPIs). Routing Director evaluates these metrics against the threshold set in the evaluation criteria. For example, you can set that the threshold for delay to be below 10 milliseconds. If such a threshold is violated an event is raised with the configured severity.

Test Agents send streams from the source device—on which you have installed the Test Agent—to a remote endpoint. The remote endpoint can be a remote Test Agent, Web servers, DNS servers, or a known device in the public cloud provider's network.

To summarize, you create a Test or a Monitor using the Routing Director GUI or REST API to configure the plug-ins, the metrics, the evaluation criteria per metric, and so on. When you run the Test or the Monitor, Test Agent starts downloading the plug-in and starts the measurement and their associated Streams.

For information on Active Assurance Terminologies, see ["Active Assurance Terminologies" on page 892](#).

For information on creating Tests and Monitors, see ["Create a Test" on page 958](#) and ["Create a Monitor" on page 1110](#).

Benefits of Using Test Agents

- You can install Test Agents on any x86 platform or as a Docker container.
- You can automate testing and data collection using Test Agents.
- You can test network performance at specific location before deploying a business critical application.

RELATED DOCUMENTATION

[About the Test Agents Page](#) | 900

About the Test Agents Page

IN THIS SECTION

- [Tasks You Can Perform](#) | 901

To access this page, click **Inventory > Active Assurance > Test Agents**.

Test Agents evaluate the quality of your network by collecting metric data for preconfigured key performance indicators (KPIs). Test Agents deployed in different locations in your network send metric data to Routing Director over management interfaces. For more information on Test Agents, see ["Test Agents Overview" on page 897](#).

You (superusers and network administrators) can view the following widgets on the Test Agents page:

- Urgent actions needed—Displays all the critical issues that require immediate attention or intervention as the issues impact the functioning of the Test Agent.
- Actions needed—Displays all the major issues that require attention as the issues impact the functioning of the Test Agent.
- Total—Displays the total number of Test Agents in the network.

- In Use—Displays the number of Test Agents that are actively running a Test or a Monitor.



NOTE: The term **Test Agent** refers to both the Test Agent Application and Test Agent Appliance. If a specific type of Test Agent is intended, it is mentioned explicitly.

Tasks You Can Perform

You can perform the following tasks on the Test Agents GUI:

- View details of Test Agents—You can view the list of all the Test Agents in your network and you can also view details of a Test Agent you have selected. To view the details of Test Agents, select a Test Agent and click **More > Details**. You can also hover over the *Test-Agent-Name* and click the **Detailed View** icon. The Test Agent Details pane appears on the right side of the page displaying the Test Agent details.

On the Test Agent Details pane, you can:

- **Copy API Request URL**—You can copy the API Request URL so that you can reuse the copied API request URL to fetch the Test Agent details.
- View information such as, name, description, tags, ID, types, online, creation, and execution status of the Test Agents.
- Edit the name, description, and tags of the Test Agent.

Click the Close (x) icon to close the pane.

- Add a Test Agent—Routing Director allows you to discover a Test Agent Application. For Routing Director to discover a Test Agent Application, you need to add a Test Agent Application by using the **Add Test Agent** button. For more information, see ["Add a Test Agent" on page 908](#).
- View details of a Test Agent—To view the details of a selected Test Agent, click a *Test-Agent-Name*. The Test Agent Details page displays the following tabs:
 - **Overview**—Displays system information and location details of the Test Agent.
 - **Interfaces**—Displays all the interfaces used by the Test Agent.
 - **Download & Run**—Displays the secret key and commands required to download, install, and run a new Test Agent Application in a Docker environment. For more information, see ["About the Test Agent-Name Page" on page 911](#).
- You can also perform the following tasks on this page:
 - Sort, resize, or re-arrange columns in a table (grid).

- Show or hide columns in the table or reset page preferences, using the vertical ellipsis menu.
- Filter the data displayed in the table—Click the filter icon (funnel) and select whether you want to show or hide advanced filters. You can then add or remove filter criteria, save criteria as a filter, apply or clear filters, and so on. The filtered results are displayed on the same page.

For more information, see ["GUI Overview" on page 6](#).

Table 254: Description of Fields on the Test Agents Page

Fields	Description
ID	The Universal Unique Identifier (UUID) of the Test Agent.
Name	The name of the Test Agent that you have specified at the time of installing a Test Agent. Based on the icon of the Test Agent, you can identify the type of Test Agent. • Blue circle icon—Indicates that the Test Agent Application is standalone and the Test Agent is not associated with any device. • Router icon with blue circle on top—Indicates that the Test Agent Application is onboarded and deployed on a device and the Test Agent is associated with the device. • Router icon—Indicates that the Test Agent is associated with a device but does not support Test Agent Application. • Orange square icon—Indicates that the Test Agent is a standalone Test Agent Appliance.
Description	The description of the Test Agent that you have specified at the time of configuring the Test Agent. You can also edit the description from the <i>Test-Agent</i> Details page.
Online	The connection status of the Test Agent. For more information, see Table 255 on page 905 .

Table 254: Description of Fields on the Test Agents Page (*Continued*)

Fields	Description
Device Online	The connection status of the device on which the Test Agent is installed. For more information, see Table 255 on page 905 .
In Use	In use indicates whether the Test Agent is actively collecting traffic measurements or not. True—Test Agent is actively collecting traffic measurements. False—Test Agent is not collecting traffic measurements.
Events	The total number of Test Agent-related events generated by Routing Director. The Events column displays the number of issues that are related to Test Agents or measurements related to Test Agents. It also displays icons representing different severity levels in the order of priority. The highest priority is marked as Critical, indicating the most urgent issues, while the lowest priority is labeled as Informational, signifying minor notifications. Click the <i>Event-Number-link</i> and the Events page appears displaying the list of Test Agent-related events generated by Routing Director. For more information, see Table 256 on page 907.
IPv4 Address	The IPv4 address of the device interfaces from which Routing Director initiates the test.
IPv6 Address	The IPv6 address of the device interfaces from which Routing Director initiates the test.
External IP (Address)	The public IP address of the Test Agent. Based on the Test Agent deployment, this IP address is a static or dynamic (using DHCP) assignment.

Table 254: Description of Fields on the Test Agents Page *(Continued)*

Fields	Description
Software version	The software version of the Test Agent.
Device ID	The UUID of the device on which the Test Agent is deployed.
Device Model	The model of the device on which the Test Agent is deployed.
Device Name	The name of the device on which the Test Agent is deployed.
Device State	The installation status of Test Agent on a device. For example, Configuring, Installing, Installed, Deconfiguring, Uninstalling, Uninstalled, Failed, or Error.
Device Requested State	The installation status that you have requested during the process of installing a Test Agent on a device. For example, if you have requested to install a Test Agent on a device, the Device State can be anything from Configuring to Error, but the Device Requested State will be Installed.
Device Version	The OS version of the device on which the Test Agent is deployed. For example, 23.2R1.15-EVO.

Table 254: Description of Fields on the Test Agents Page (*Continued*)

Fields	Description
Tags	The tags configured for the Test Agent in the <i>key:value</i> format. By default, only two tags are displayed in the column. To view more than two tags, click + Number more. A tag is a key-value pair in which the key signifies a category for which you configure a value. The value is an identifier for the category. Examples of key-value pairs are device and device name (edgedevice:acx7000), site and site name (site:bangalore). You can configure tags from the <i>Test-Agent Details</i> page.

Table 255: Connection status of Test Agents and Devices

Online (Indicates the connection status of the Test Agent)	Device Online (Indicates the connection status of the device)	Status (Describes the type of the Test Agent)
Green active Wi-Fi icon	Green active Wi-Fi icon	Indicates that the Test Agent is associated with a device and the Test Agent Application is running. Both the device and Test Agent are online.
Red inactive Wi-Fi icon	Green active Wi-Fi icon	Indicates that the Test Agent is associated with a device and the Test Agent Application is running. The device is online but the Test Agent is offline.

Table 255: Connection status of Test Agents and Devices *(Continued)*

Online (Indicates the connection status of the Test Agent)	Device Online (Indicates the connection status of the device)	Status (Describes the type of the Test Agent)
Red inactive Wi-Fi icon	Red inactive Wi-Fi icon	Indicates that the Test Agent is associated with a device and the Test Agent application is running. Both the device and the Test Agent is offline.
Green active Wi-Fi icon	N/A	Indicates that the Test Agent Application is running independently and is not associated with any device. The Test Agent is online.
N/A	Green active Wi-Fi icon	Indicates that the Test Agent is associated with a device but has no Test Agent application running. The device is online.
Red inactive Wi-Fi icon	N/A	Indicates that the Test Agent Application is running independently and is not associated with any device. The Test Agent is offline.
N/A	Red inactive Wi-Fi icon	Indicates that the Test Agent is associated with a device but has no Test Agent Application running. The device is offline.

Table 256: Description of fields on the Events page

Field	Description
Severity	The severity level of the event that is raised by Test Agents. The following are the severity levels: Critical—Indicates that the event is critical and the issue needs immediate attention. Warning—Indicates that the issue that has caused the event needs to be fixed but does not require immediate attention. Info—Indicates the progress of a task and provides information on the event. The informational event does not require immediate attention. Error—Indicates that the issue that has caused the event needs to be fixed and requires immediate attention and troubleshooting.
Description	The description that you specified when you configured the evaluation criteria for a Test.
Raise Time	The date and time when the event was generated. The timestamp is displayed in the following format: Month DD, YYYY, HH:MM:SS AM/PM. The date and time is displayed according to the Raise delay you specified when you configured the evaluation criteria for a Test. For example, Sep 5, 2024, 4:29:52 PM.
Clear Time	The date and time when the event was cleared. The timestamp is displayed in the following format: Month DD, YYYY, HH:MM:SS AM/PM. The date and time is displayed according to the Clear delay you specified when you configured the evaluation criteria for a Test. For example, Sep 5, 2024, 4:29:52 PM.

Table 256: Description of fields on the Events page *(Continued)*

Field	Description
Subject	Click the Show link to view the details of event generated in JSON format. It displays Test Agent ID, Test ID, Task ID, Stream ID, and so on.
Data	Click the Show link to view the details of the evaluation criteria associated with the event generated.

RELATED DOCUMENTATION

| [Tests and Monitors Overview](#) | 944

Add a Test Agent

IN THIS SECTION

- [Add a Test Agent Application:](#) | 909
- [Add a Test Agent Appliance](#) | 910

A Test Agent is a software that is installed on your network device, and it acts as a measurement point in your network. These Test Agents are deployed at strategic locations in your network to evaluate the quality of your network by collecting metric data for preconfigured key performance indicators (KPIs). For more information, see "[Test Agents Overview](#)" on page 897.

There are two types of Test Agents:

- **Test Agent Application**—A lightweight Test Agent that runs within a docker container on an existing operating system. This type of Test Agent does not support interface configuration and certain advanced use cases. You can deploy a Test Agent Application natively in an x86 platform or as a docker container.

- **Test Agent Appliance**—A full-fledged Test Agent with built-in operating system that allows full control over network configuration and supports advanced functionalities. You can deploy a Test Agent Appliance as a virtual machine or a custom x86 hardware.

Add a Test Agent Application:

To add a Test Agent Application:

1. Navigate to the Test Agents (**Inventory > Active Assurance > Test Agents**) page.
2. Click the **Add Test Agent** button.

The Add Test Agent page appears.



NOTE: Fields marked with an asterisk (*) are mandatory.

3. Select Test Agent Application as type.
4. Specify the Name, Tags, and Description for the Test Agent.
5. Click **Create**.

A message confirms that Routing Director has successfully added the Test Agent, and you are redirected to the Test Agent Details (**Inventory > Active Assurance > Test Agents > Test Agent Details**) page.

On the Download & Run tab, view the version of the Test Agent Application. By default, the version is prepopulated with the latest version of the Test Agent Application.

6. On the **Download & Run tab**, copy the secret key generated by Routing Director by using the **Copy** button next to the Secret text box. You can use this secret key in the docker environment for authenticating the Test Agent Application.

The secret key is displayed only once. If you return to the **Download, Install & Run tab** later, you need to regenerate the secret key by using the **Regenerate** button.

7. Copy the commands by using the **Copy** button to download, install, and run the Test Agent Application in the Docker environment.
8. Log in to Docker environment and execute the copied secret key and commands in the Docker environment, which allows the Test Agent Application to establish a connection with Routing Director. For more information, see ["Install Test Agent Application " on page 929](#).

Now, Routing Director discovers the Test Agent Application, and the status of the Test Agent Application becomes online.

Add a Test Agent Appliance

To add a Test Agent Appliance:

1. Navigate to the Test Agents (**Inventory > Active Assurance > Test Agents**) page.
2. Click the **Add Test Agent** button.

The Add Test Agent page appears.

3. Select Test Agent Appliance as the type of Test Agent.
4. Specify the Name, Tags, and Description for the Test Agent.
5. Click **Create**.

A message confirms that Routing Director has successfully added the Test Agent Appliance, and you are redirected to the Test Agent Details (**Inventory > Active Assurance > Test Agents > Test Agent Details**) page.

6. On the **Download & Run tab**, copy the secret key generated by Routing Director by using the **Copy** button next to the Secret text box. You can use this secret key in the virtualization platforms to authenticate the Test Agent Appliance.

The secret key is displayed only once. If you return to the Download, Install & Run tab later, you need to regenerate the secret key by using the **Regenerate** button.

7. Copy the **Cloud-init configuration** by using the **Copy** button at the bottom-right of the Cloud-init configuration text box and use this configuration on various virtualization platforms. In addition, you can enable the **Include examples** checkbox to view a complete template with a reference of possible Cloud-init configurations.
8. Follow the installation procedures dedicated to any virtualization platform and configure the cloud-init metadata to setup Test Agent Appliance, which allows Routing Director to discover the Test Agent Appliance, and the status of the Test Agent Appliance becomes online. For more information, see ["Install Test Agent Appliance" on page 934](#).

RELATED DOCUMENTATION

| [Active Assurance Workflow](#) | 891

About the *Test Agent-Name* Page

IN THIS SECTION

- [Tasks You Can Perform | 911](#)
- [Overview Tab | 914](#)
- [Interfaces Tab | 918](#)
- [Download & Run Tab | 927](#)

To access the *Test-Agent-Name* page:

1. Select **Inventory > Active Assurance > Test Agents**.

The Test Agents page appears.

2. Click a *Test-Agent-Name*.

The Test-Agent-Name page appears displaying the Overview, Interfaces, and Download & Run tabs.

Tasks You Can Perform

You can perform the following tasks from this page:

- View the connection (Online and Device Online) status of the configured Test Agent.
- View the total number of Test Agents that are in-use.
- View the total number of events generated by the Test Agent.
- View the installation status of the Test Agent on the device.
- View the Test Agent Gateway Hostname (IPv4 and IPv6) addresses.
- Commit configuration changes—You can apply all the configuration changes you have made for the global (management and NTP interfaces) and network interfaces by using the **Commit Changes** button. To apply the configuration changes made in the **Configuration** accordion and **Interfaces** tab:

1. Click the **Commit Changes** button.

A Create New Commit page appears.

You can also enable the toggle button to automatically rollback if the commit fails.

2. Click **Commit** to apply configuration.

You are taken to the Commit page where you can view the commit status, settings, details, and full interface configuration details. An alert appears displaying the progress of your commit on the top of the Test-Agent-Name page.

- Discard configuration changes—You can discard all the configuration changes you have made for the global (management and NTP interfaces) and network interfaces by using the **Discard Changes** button. To discard the configuration changes:

- Click the **Discard Changes** button.

A Discard all pending configuration changes page appears asking you to confirm the discard operation on all pending configuration changes.

- Click **Discard Changes**.

A message is displayed confirming that you have successfully discarded the changes.

- View configuration history—You can view the history of configuration changes that were applied to Test Agent Appliance. Click the **Commit history** button and a Test Agent Config Commit History page appears with a list of configuration changes displaying the status, status message, the user address of the person, timestamp, and so on. From this page, you can also apply or discard configuration changes.
- Upgrade a Test Agent Appliance—You can upgrade a Test Agent Appliance to its latest software version.

To upgrade a Test Agent Appliance:

1. Click **More > Upgrade**.

A *Test-Agent-Appliance* Upgrade page appears asking you for a confirmation on upgrade.

2. Click **Upgrade** to confirm the upgrade operation.

The Test Agent Appliance you selected upgrades to the latest version, and you can track the progress on the Software accordion of the Overview tab (**Inventory > Active Assurance > Test Agents > Test Agent Details page > Overview**).



WARNING: When you upgrade a Test Agent Appliance, the Test Agent Appliance goes offline and all active measurements related to Test Agent Appliances are disrupted.

- Reboot a Test Agent—You can reboot Test Agent Appliance.

To reboot a Test Agent Appliance:

1. Click **More > Reboot**.

A Reboot Test Agent Appliance page appears asking you for a confirmation the reboot operation.

2. Click **Yes** to confirm the reboot operation.



WARNING: When you reboot a Test Agent Appliance, the Test Agent Appliance goes offline and all active measurements related to Test Agent Appliances are disrupted.

- Delete a Test Agent—To delete a Test Agent:

1. Click **More > Delete**.

2. Click **Yes** to confirm the delete operation.

A message is displayed confirming that the Test Agent is deleted successfully.



NOTE: If you delete a Test Agent, the Tests that are associated with this Test Agent fail and the Monitors that are associated with this Test Agent stop running.

- View the following tabs:

- Overview Tab—You can view the connection status, the location of the Test Agent, system information to troubleshoot issues, statistics of the interfaces of the Test Agent, Test Agent upgradation details, and you can also configure the global (management and NTP interface) and network interfaces of Test Agent Appliances. For more information, see ["Overview Tab" on page 914](#).
- Interfaces Tab—You can view the interface details relevant to the Test Agent such as IPv4 and IPv6 addresses, interface names, and their descriptions. In addition, you can configure interfaces details and state. For more information, see ["Interfaces Tab" on page 918](#).



NOTE: If the Test Agent is offline, Routing Director displays the details of interfaces, system information, and location status that were retrieved when the Test Agent was online.

- Download & Run Tab—You can view the Test Agent installation instructions.
- For Test Agent Application—You can generate and copy the secret key as well as the commands to download, install, and run the Test Agent Application in a Docker environment.

- For Test Agent Appliance—You can view and copy the cloud-init configuration to set up Test Agent Appliance on various virtualization platform. For more information, see "[Download & Run Tab](#)" on page 927.

Overview Tab

You can access the **Overview** tab on the *Test-Agent-Name* (**Inventory > Active Assurance > Test Agents > Test-Agent-Name**) page.

In the Overview tab, you can:

- View the identity and location, the system information of software environment, and the interface statistics of the Test Agent. In addition, you can view the configuration and software upgradation details of Test Agent Appliances.

By default, all accordions are expanded in the **Overview** tab. You can also click to expand and minimize each accordion to view the following information:

- **Identity & Location**—Click and expand the **Identity & Location** accordion. In this accordion, you can view the following information:
 - **Metadata**—View and edit the name, description, and tags of the Test Agent.
 - **Location**—View the site details of a Test Agent and assign the Test Agents to a site. To assign a Test Agent to a site:
 1. Click the **Edit** (pencil) icon near the site name.

The Assign Test Agents to Site page appears displaying the details described in [Table 259 on page 918](#).
 2. Select a site that you want to assign from the available list.
 3. Click **OK**.

A message confirms that Routing Director has successfully updated the site.



NOTE: You can assign a Test Agent to a site only if the Test Agent is not installed on any devices. For the Test Agent installed on a device, the site is pre-populated with the location of that device. To update the site information of a device, see "[About the Sites Page](#)" on page 88.

- **Test Agent**—View the ID, created and updated details, last connect time, and last disconnect time of the Test Agent. You can also copy the ID of the Test Agent.

- **Device**—View the ID, name, model, MAC address, and installation state of the device on which you have associated the Test Agent. You can also copy the ID of the device.
- **Configuration**—Click and expand the **Configuration** accordion to view the following information on global interface configuration:
 - **Management interface configuration**—View the list of network interfaces used for managing Test Agent Appliance. You can select the desired interface from the list of interfaces.
 - **NTP interface configuration**—Select the network interface to synchronize the time of Test Agent Appliance with NTP server. You can:
 - Enable the **Same as Management Interface** toggle button if you want to use the management interface for NTP communication.
 - Enable the **Enable IPv6 for NTP** toggle button to support IPv6 for NTP synchronization.
 - Specify IP address of NTP servers for time synchronization.
 - **NTP Stats**—View status of time synchronization. The statuses include reachability, poll interval, last received time, and so on.
- **Software**—Click and expand the **Software** accordion to view the following information on software upgradation of Test Agent Appliance:
 - **Device**—View the type and vendor of the device. By default, the type is Appliance.
 - **Versions**—View the software version details such as current version, upgradation status of Test Agent Appliance. In addition, you can also upgrade Test Agent Appliance to the latest available version.
- **Interfaces**—Click and expand the **Interfaces** accordion to display the following information on the interfaces:
 - Number of interfaces.
 - Number of interfaces that are online.
 - Number of interfaces that are offline.
- **System Information**—Click and expand the **System Information** accordion to view the information related a Test Agent as described in the [Table 258 on page 916](#).

Table 257: Fields on the NTP Stats Section

Field	Description
Status	Displays the current synchronization status of the NTP server. For example, reachable, unreachable.
Address	Displays the IP address of the NTP server.
Stratum	Displays the number of hierarchical hops between server and the reference clock. Lower the stratum number, higher the accuracy. For example, Stratus 2 indicates that the server is 2 hops away from the clock and accuracy is low compared to stratus 1
Poll interval	Displays the interval between consecutive requests that Test Agent Appliance sends to the NTP server. Low poll interval indicates connection issues and long poll interval indicates that the synchronization is stable.
Reach	Displays the number of recent successful NTP requests.
Last RX	Displays the number of polling intervals since the last successful NTP response. Low RX value indicates that Test Agent recently synced with NTP server and high RX value indicates network or server issues.

Table 258: System Information Fields of a Test Agent

Field	Description
Agent Version	Displays the software version of the Test Agent. For example, 4.3.0.18

Table 258: System Information Fields of a Test Agent (*Continued*)

Field	Description
Hostname	Displays the hostname of the device on which the Test Agent is installed.
Kernel Version	Displays the kernel version of the Test Agent software.
Architecture	Displays the software architecture and its operation mode in the device. For example, x86_64.
Bios Version	Displays the BIOS version of host device.
System Manufacturer	Displays the manufacturer of the host device. For example, Juniper.
System Product Name	Displays the product name of the host device. For example, ACX7024
System Version	Displays the version number of the OS running on the host device. For example, 23.2R1.15-EVO.
Processor Manufacturer	Displays the manufacturer of the processor available on the host device.
Processor Version	Displays the version number of the processor available on the host device.
Management Interface	Displays the name of the management interfaces that the Test Agent uses to communicate with Routing Director. For example, re0:mgmt-0.0.

Table 258: System Information Fields of a Test Agent (*Continued*)

Field	Description
Management Address	Displays the IP address of the management interface that the Test Agent uses to communicate with Routing Director. For example, 10.48.54.198.

Table 259: Fields on the Assign Test Agents to Site Page

Fields	Description
Name	Displays the name of the site.
Country	Displays the country where the site is located.
Timezone	Displays the time zone of the site.
Address	Displays the address of the site.
Notes	Displays additional information about the site.

Interfaces Tab

You can access the **Interfaces** tab on the *Test-Agent-Name* (**Active Assurance** > **Test Agents** > *Test-Agent-Name*) page.

In the Interfaces tab, you can:

- View details of interfaces—You can view the general information such as the name and the IP addresses of all the interfaces associated with Test Agent as described in the [Table 260 on page 921](#).
- Hide Orphaned interfaces—The orphaned interfaces are network interfaces that are not associated with any active configurations or the interfaces that are newly created but not yet configured within Routing Director. You can select the **Hide Orphaned Interfaces** checkbox to exclude orphaned interfaces from the Interfaces list. By selecting the checkbox, you can view only the active interfaces.
- Create an interface—You can create a new interface for network expansion and traffic isolation by using the **Add new Interface** button on the top-right corner of the page.

To create an interface:

1. Click the **Add new Interface** button.

A Create New Interface page is displayed.

2. Configure the interface details as described in [Table 261 on page 923](#).



NOTE: Fields marked with an asterisk (*) are mandatory.

3. Click **Create**.

A message is displayed confirming that you have successfully created an interface, and an *interface-name* page appears where you can configure fields as described in [Table 262 on page 924](#).

- Configure an interface—You can either configure an interface while creating an interface by using the **Add new Interface** button, or by clicking an *interface-name* from the Interfaces tab. You can specify the interfaces details and view interfaces state as described in [Table 262 on page 924](#).

You can also set the interface as management interface and set the interface as NTP source for Test Agent Appliances.



NOTE: You can only configure interfaces for Test Agent Appliances. For Test Agent Applications, you can only view the details related to the state of interfaces.

- Delete an interface—You can delete a network interfaces based on their association status with Test Agents. You can delete interfaces such as VLAN, bridge, and physical types under specific conditions.

For non-orphaned interfaces on Test Agents, mark an interface for deletion by using the **Delete** functionality and then commit the changes.

For orphaned interfaces on any Test Agents, delete the interfaces directly by using the **Delete** functionality. If needed, you can also unmark an interface for deletion by using the **Undelete** functionality, which restores the interface before the configuration is committed.



NOTE:

- Currently, only VLAN interfaces are supported for Test Agent Appliances.
- If an interface is orphan, you can delete the selected interface directly without committing the change.

To delete a Test Agent Application interface:

1. Select the interface that you want to delete.

2. Click **More** and select **Delete**.

A warning page is displayed asking you to confirm the delete operation.

3. Click **Yes**.

The interface you have selected is successfully deleted.

To delete a Test Agent Appliance interface:

1. Select the VLAN interface that you want to delete.

2. Click **More** and select **Delete**.

A warning page is displayed asking you to confirm the delete operation.

3. Click **Yes**.

4. Click the **Commit Changes** button to apply the delete configuration.

The interface you have selected is successfully deleted.

To undelete a Test Agent interface:

1. Select a VLAN or an orphan interface that you want to restore.

2. Click **More** and select **Undelete** option.

A message is displayed asking you to confirm the undelete operation.

3. Click **OK**.

The interface you have selected is successfully restored.



NOTE: You can restore (undelete) a VLAN interface or orphaned interface only if it was marked for deletion and the changes have not been committed yet. If the changes are already committed, the interface cannot be restored.



TROUBLE: After you undelete an interface, commit operations for the affected Test Agent fails regardless of the interface involved.

To resolve this issue, you can delete the orphan VLAN interfaces. This will restore the commit functionality for the affected Test Agent.

- You can also perform the following tasks on this page:

- Sort, resize, or re-arrange columns in a table (grid).
- Show or hide columns in the table or reset page preferences, using the vertical ellipsis menu.
- Filter the data displayed in the table—Click the filter icon (funnel) and select whether you want to show or hide advanced filters. You can then add or remove filter criteria, save criteria as a filter, apply or clear filters, and so on. The filtered results are displayed on the same page.

Table 260: Fields on the Interface Tab

Fields	Description
Interface Name	Displays the interface associated with the Test Agent Appliance.
Namespace	Displays the network namespace on which the interface is associated.
Link	Displays the status of the links associated with the interface.
MAC	Displays the unique MAC addresses associated with the interfaces.
IPs	Displays the list the IP addresses associated with the interface. Displays both IPv4 and IPv6 addresses.
Type	Displays the type of interface. By default, the configuration type is pre-populated as Physical.
Management IP	Displays the management IP address used for managing the device remotely.
Assigned IPV4 Addresses	Displays the IPv4 addresses currently assigned to the interface.
Assigned IPV6 Addresses	Displays the IPv6 addresses currently assigned to the interface.
Running Config - Static IPV4/IPV46 addresses	Displays the manually configured IPv4/IPv6 addresses in the running configuration.

Table 260: Fields on the Interface Tab *(Continued)*

Fields	Description
Running Config - Static Nameservers	Displays the manually configured DNS nameservers in the running configuration.
Running Config - Static Routes	Displays the static routes configured in the running configuration.
Running Config - DHCP Vendor ID (IPv4)	Displays the DHCP Vendor ID for IPv4 addresses in the running configuration.
Running Config - Static IPv6 Addresses	Displays the manually configured IPv6 addresses in the running configuration.
Running Config - Static IPv6 Nameservers	Displays the manually configured IPv6 nameservers in the running configuration.
Running Config - Static IPv6 Routes	Displays the manually configured IPv6 routes configured in the running configuration.
Running Config - DHCP Vendor ID (IPv6)	Displays the DHCP Vendor ID for IPv6 addresses in the running configuration.
Running Config - SLAAC Nameservers	Displays the nameservers received via Stateless Address Autoconfiguration (SLAAC) in the running configuration.
Running Config - First Received Time	Displays the time when the first configuration data was received by the interface.
Running Config - ID	Displays the unique identifier for the running configuration of the interface.
Pending Config - MTU Size	Displays the MTU size pending for the interface configuration.
Pending Config - MAC Address	Displays the MAC address that are yet to be assigned to the interface.

Table 260: Fields on the Interface Tab *(Continued)*

Fields	Description
IPv4 External IP Type	Displays the type of external IPv4 address associated with the interface.
IPv4 External IP Address	Displays the external IPv4 address assigned to the interface.
IPv6 External IP Type	Displays the type of external IPv6 address associated with the interface.
IPv6 External IP Address	Displays the external IPv6 address assigned to the interface.
Pending Config - Static Routes	Displays the static routes that are yet to be configured for the interface.
Created by	Displays the email address of the user who created the interface.
Create time	Displays the date and time when the interface was created.
Updated by	Displays the email address of the user who updated the interface details.

Table 261: Fields on the Create New Interface Page

Fields	Description
Type	Select the type of the interface. Currently, only VLAN interfaces are supported for Test Agent Appliances.
Parent*	Select the physical interface on which you want to create the VLAN.
VLAN ID*	Specify the unique VLAN identifier (ID) for the interface.

Table 261: Fields on the Create New Interface Page *(Continued)*

Fields	Description
Namespace Isolation	Namespace isolation type determines how you want to assign the VRF and namespace to the VLAN interface. • Enable the Same As Parent (VRF) toggle button for the VLAN interface to share the parent namespace. • Enable Isolated for the VLAN interface to get a new interface based on the interface name that you are configuring.
Name	The name field is automatically populated with the Parent and VLAN ID in <i>parent.vlanID</i> format. For example, <i>eth2.3</i>.
Namespace	Specify the namespace of the VLAN interface. • If you select Same As Parent (VRF), the namespace inherits the name of the parent (physical) interface. • If you select Isolated, the namespace automatically populates the value you specified in the Name field.

Table 262: Fields on the *Interface-Name* (namespace: *Namespace-Name*) Page

Fields	Description
Type	By default, the Configuration type is pre-populated as Physical.
Type (IPv4)	Select the type of IPv4 address you want to assign to the interface. • None—To not assign IPv4 address to the interface. • DHCP—To dynamically assign IPv4 address to the interface. • Static—To manually assign IPv4 address to the interface.

Table 262: Fields on the *Interface-Name* (namespace: *Namespace-Name*) Page (Continued)

Fields	Description
Type (IPv6)	Select the type of IPv6 address you want to assign to the interface. • None—To not assign IPv4 address to the interface. • DHCP—To dynamically assign IPv4 address to the interface. • SLAAC—To use Stateless Address Autoconfiguration. • Static—To manually assign IPv4 address to the interface.
Advanced	
Speed	Select the speed of the network interface. For example, 10Mbps Half Duplex, 10Mbps Full Duplex.
MTU	Specify the maximum packet size that can be transmitted over the network. If not specified, the hardware MTU size is used.
Override MTU	Enable the Override MTU toggle button to manually set a custom MTU instead of default value.
Mac	Specify the Media Access Control (MAC) address of the interface.
Spoof MAC address	Enable the Spoof MAC address toggle button to manually set a MAC address instead of default address assigned by hardware.
State	
Link Carrier	Indicates the carrier link status of interface: • True—Signal is detected. • False—Signal is not detected.

Table 262: Fields on the *Interface-Name* (namespace: *Namespace-Name*) Page (Continued)

Fields	Description
Link Admin	Indicates whether the interface is currently active or inactive. • Up—Interface is active. • Down—Interface is inactive.
IPv4 Addresses	Indicates the IPv4 addresses assigned to the interface.
IPv6 Addresses	Indicates the IPv6 addresses assigned to the interface.
MTU	Indicates the maximum size of packets that the interface can transmit.
MAC	Indicates the MAC address of the interface.
Speed	Indicates the link speed of the selected interface in Mbps.
Duplex Mode	Indicates whether the interface operates in full-duplex (send and received simultaneously) or half-duplex (one direction at a time).
Hardware Timestamping	Indicates whether the interface supports packet timestamping
External IP	Indicates the external-facing IP address assigned to the interface.
PCI Vendor	Indicates the name of the vendor of network interface card.
PCI Device	Indicates the model identifier of network interface card.
NIC Driver Name	Indicates the name of the driver that manages the interface.
NIC Driver Version	Indicates the version of the driver that manages the interface.

Download & Run Tab

- **For Test Agent Application:**

You can access the **Download & Run** tab on the Test Agent Details (**Active Assurance > Test Agents > Test-Agent-Application-Name**) page.

On the Download & Run tab, you can:

- View the following information required to configure the Test Agent Application:
 - **Version**—You can view the version of the Test Agent Application. By default, the version number is pre-populated with the latest version of the Test Agent Application.
 - **Secret**—You can view, copy, and regenerate the unique secret key required to authenticate and install the Test Agent Application in the Docker environment.
 - **Command to Download, Run & Install**—You can view and copy the commands required to download, install, and run the Test Agent Application in the Docker environment. The commands are added to the **Command to Download, Install & Run** text box based on the check boxes you select. The available check boxes are:
 - **Include download instruction**—Displays command required to download the Test Agent Application.
 - **Include run instruction**—Displays command required to run the Test Agent Application.
 - **Include debug logs**—Displays command required to generate the debug logs while running and installing the Test Agent Application.
 - **Ignore SSL certificates**—Displays command required to ignore the SSL certificate verification in the Docker environment.
- Regenerate the secret key—You can regenerate the secret key by clicking the **Regenerate** button. Every time you regenerate, a new authentication secret is generated.

By using the Secret key, you can reuse an existing Test Agent registered with Routing Director. Reusing the secret key allows the Test Agent to function from a different location while remaining a part of active measurements. To be a part of active measurements, you must use the same container identity.

If you regenerate a secret key, you must reconfigure all linked measurements and Test templates. Also, you cannot regenerate a secret if the Test Agent associated is online and connected.



NOTE: A secret key is a unique key used to authenticate the Test Agent Application in a Docker environment. The secret key allows only the authorized Test Agent

Applications to communicate with Routing Director and prevents unauthorized access. The secret key is displayed only once. You cannot view the same secret key if you return to the Download, Install & Run Tab later.

As you cannot view the secret key more than once, you need to use the Regenerate button and generate a new secret key.

- **Copy the secret key**—You can copy the secret by using the **Copy** button. You can copy them to clipboard and execute the key in the Docker environment while installation of the Test Agent Application.
- **Copy the commands**—You can copy the commands to download, install, and run the Test Agent Application by using the **Copy** button. You can copy them to clipboard and execute the commands in the Docker environment.

When you execute the commands from the Docker environment, the Test Agent Application is downloaded and installed into a Docker environment. The secret key authenticates the Test Agent Application and establishes a secure connection with Routing Director. Routing Director now discovers the Test Agent Application and displays the status as Online.

- **For Test Agent Appliance:**

You can access the **Download & Run** tab on the Test Agent Details (**Active Assurance > Test Agents > Test-Agent-Appliance-Name**) page.

On the Download & Run tab, you can:

- View the following information required to configure the Test Agent Appliance:
 - **Secret**—You can view, copy, and regenerate the unique secret key required to authenticate and install the Test Agent Appliance.
 - **Command to Download, Run & Install**—You can view and copy the Cloud-init configuration used to setup Test Agent Appliance on various virtualization platforms. The commands are updated on the **Cloud-init configuration** text box based on the check boxes you select. The available check boxes are:
 - **Include examples**—Displays references of various possible ways to customize the Test Agent Appliance.
 - **Ignore SSL certificates**—Displays command required to ignore the SSL certificate verification of Routing Director.
- **Regenerate the secret key**—You can regenerate the secret key by clicking the **Regenerate** button. Every time you regenerate, a new authentication secret is generated.



WARNING: Regenerating the secret invalidates the previous secret. Any Test Agent configured by using the previous secret loses its connection with Routing Director.



NOTE: A secret key is a unique key used to authenticate the Test Agent Appliance. The secret key allows only the authorized Test Agent Appliances to communicate with Routing Director and prevents unauthorized access. The secret key is displayed only once. You cannot view the same secret key if you return to the Download, Install & Run Tab later.

As you cannot view the secret key more than once, you need to use the Regenerate button and generate a new secret key.

- Copy the secret key and cloud-init configuration—You can copy the secret by using the **Copy** button beside the Secret text box. You can copy them to clipboard and execute the key in the virtualization platform while installation of Test Agent Appliance.

You can copy the cloud-init configuration used to setup Test Agent Appliance by using the **Copy** button on the bottom-right of the text box. When you use the cloud-init configuration on any virtualization platform, you register the Test Agent Appliance with Routing Director. The secret key authenticates the Test Agent Application and establishes a secure connection with Routing Director. Routing Director now discovers the Test Agent Appliance and displays the status as Online.

Install Test Agent Application

SUMMARY

The installation topic describes procedures to register the Test Agent using the official container image, build a test agent image, and deploy a native Test Agent application on an x86 machine.

IN THIS SECTION

- [Deploy a Test Agent Using Container Image | 930](#)
- [Build a Test Agent Container Image | 931](#)
- [Install a Test Agent as a Native Application in Linux | 932](#)

The Test Agent application is a container that you deploy natively in an x86 platform or as a docker container, or as a container in a router running Junos Evolved OS. In this installation procedure, we will explore the steps to deploy Test Agents as a container.

Before You Begin

Before you begin the Test Agent installation, you must:

- Ensure to install [Docker Engine](#) in your system that runs the latest Ubuntu version.
- Ensure that no firewalls are blocking the connection to Routing Director on TCP port 6800.



WARNING: We do not recommend installing Test Agents using Junos CLI.

Deploy a Test Agent Using Container Image

To deploy a Test Agent using the container image:

1. Download the latest image of the Test Agent. You can download the Test Agent software (Test Agent Application (amd64)) from the [Software Download](#) page or from the Docker Hub (netrounds/test-agent-application).
2. Register the Test Agent using the register command. This connects the Test Agent with Routing Director. When you run the register command, specify one of the following:
 - Specify the email address and password that is used at the time of Routing Director installation.
 - Specify the organization's UUID. To find your organization's UUID, go to **Administration > Settings** in Routing Director and then **copy** the organization's UUID.
 - Specify the hostname for Active Assurance Test Agent gateway IP address.

```
user@hostname:~# docker run --rm \

    -v $(pwd):/config \
    -v /var/run/netns:/var/run/netns \
    <Test Agent docker image> \

    register --config /config/agent.conf -A\
    --host <Host name> \
    --org <Organization UUID> \
    --email <Control Center user email> \
    --password <Control Center user password> \
    --name <Test Agent name>
```

3. Start a Test Agent using the run command:

```
user@host:~# docker run --network=host --cap-add=NET_ADMIN --device=/dev/net/tun -d \
    --privileged -v $(pwd):/config -v /var/run/netns:/var/run/netns \
    --log-opt max-size=10m --log-opt max-file=2 <Test Agent docker image> \
    --config /config/agent.conf -A [-T]
```

Build a Test Agent Container Image

Use this procedure to build your own container image. To build a custom Test Agent container image:

1. Download the latest image of the Test Agent. You can download the Test Agent software (Test Agent Application (amd64)) from the [Software Download](#) page.
2. Unpack the tar file of the Test Agent software by using the following command:

```
user@host:~# tar -xvzf paa-test-agent-application_<version>_architecture.tar.gz
    #Remove the version from the directory name
user@host:~$ mv paa-test-agent-application_<version> paa-test-agent-application/
```

3. Create a docker file in the directory where you unpacked the Test Agent software. The following docker file copies the Test Agent software to the **opt/paa-test-agent-application** location and configures the entry point to run the Test Agent software.

```
FROM debian:buster-slim

RUN apt-get update && \
    apt-get install -y ca-certificates iproute2 socat && \
    rm -rf /var/lib/apt/lists/*

COPY paa-test-agent-application/ /opt/paa-test-agent-application/
ENTRYPOINT ["/opt/paa-test-agent-application/paa-test-agent-application"]
```

4. Build the docker image by using the docker build command. Replace *mycustom* in the following command with any name. The name you specify here is the name of the docker that must used going forward.

```
user@host:~$ docker build -t mycustom/paa-test-agent-application
```

5. Register the container that you created by using the following command:

```
user@hostname:~# docker run --rm \

    -v $(pwd):/config \
    -v /var/run/netns:/var/run/netns \
    <Test Agent docker image> \

    register --config /config/agent.conf -A \
    --host <Host name> \
    --org <Organization UUID> \
    --email <Control Center user email> \
    --password <Control Center user password> \
    --name <Test Agent name>
```

6. Start the Test Agent.

```
user@host:~# docker run --network=host --cap-add=NET_ADMIN --device=/dev/net/tun -d \
    --privileged -v $(pwd):/config -v /var/run/netns:/var/run/netns \
    --log-opt max-size=10m --log-opt max-file=2 <Test Agent docker image> \

    --config /config/agent.conf -A [-T]
```



TIP: The `--privileged -v` and `/var/run/netns:/var/run/netns` arguments enable Test Agents to use the network namespaces to route the traffic independently.

The `A` option enables Test Agent to detect all available namespaces and their interfaces.

The `log-opt` option is used to limit the size and the number of log files generated by Test Agent.

The `docker logs <container ID>` command enables you to inspect the container log files.

Install a Test Agent as a Native Application in Linux

You can install a Test Agent application natively on an x86 machine.

1. Download the latest image of Test Agent. You can download the Test Agent software (Test Agent Application (amd64)) from the [Software Download](#) page.

2. Unpack the tar file and install the application by using the following commands.

```

user@hostname:~$ tar -xvzf paa-test-agent-application_version_architecture.tar.gz

# Move the contents of the directory to a more permanent location

user@hostname:~$ sudo mv paa-test-agent-application_version /opt/paa-test-agent-application

# Set root as the user and group of the new directory

user@hostname:~$ sudo chown -R root:root /opt/paa-test-agent-application

```

3. Register the Test Agent by doing one of the following.

- a. Register the Test Agent by using email and password.

```

user@hostname:~$ sudo /opt/paa-test-agent-application/paa-test-agent-application register \
--config /etc/paa-test-agent-application.conf \
--host <Host name> \
--org <Organization UUID> \
--email <Control Center user email> \
--password <Control Center user password> \
--name <Test Agent name>

```

- b. Register the Test Agent by using the API token.

```

user@hostname:~$ sudo /opt/paa-test-agent-application/paa-test-agent-application register \
--config /etc/paa-test-agent-application.conf \
--host <Host name> \
--org <Organization UUID> \
--api-token <API token for user in Control Center> \
--name <Test Agent name>

```

4. Start the Test Agent application.

```
user@hostname:~$ /opt/paa-test-agent-application/paa-test-agent-application run\
--config /etc/paa-test-agent-application.conf -A
```



TIP: The `A` option enables Test Agent to detect all available namespaces and their interfaces.

SEE ALSO

[Active Assurance Workflow | 891](#)

Install Test Agent Appliance

SUMMARY

The installation topic describes procedures to install the Test Agent Appliance as a bootable image on a custom x86 hardware, install the Test Agent Appliance on a Proxmox server and register a Test Agent Appliance from the local console.

IN THIS SECTION

- [Before You Begin | 935](#)
- [Install the Test Agent Appliance on a custom x86 hardware | 935](#)
- [Install Test Agent Appliance on Proxmox | 940](#)
- [Format of cloud-config metadata | 942](#)



NOTE: Test Agent Appliance is a beta feature in this release.

A Test Agent Appliance is a full-fledged Test Agent designed to test network performance, which is integrated with an optimized Debian Linux operating system. The Test Agent Appliance can be packaged and delivered in a number of ways:

- **Dedicated Test Agent**—You can download a Test Agent Appliance Software image and install on a custom x86 hardware.

- Test Agent Virtualized Network Function (TA VNF)—You can upload a Test Agent Software image to a virtualization platform and run on a virtual machine (VM) running on a hypervisor.

Before You Begin

Before you begin the Test Agent Appliance installation:

- Ensure that no firewalls are blocking the connection to Routing Director on TCP port 6800.
- Ensure that the platform on which you are installing Test Agent Appliance has a minimum of 4.1 GiB storage space.

Install the Test Agent Appliance on a custom x86 hardware

You can boot a Test Agent Appliance directly from a USB flash drive. You can use the USB flash drive-based deployment as an alternative to the virtual machine-based deployment.

To install a Test Agent Appliance on a custom x86 hardware, you must:

1. Create a bootable USB flash drive by using command-line tool such as dd or applications such as Etcher or Win32 Disk Image. For more information, see ["Create a bootable USB flash drive" on page 935](#).
2. Boot the custom x86 hardware from USB flash drive by using the BIOS boot manager. For more information, see ["Boot the custom x86 hardware from USB flash drive" on page 936](#).

Create a bootable USB flash drive

You can create a bootable USB flash drive in one of the following ways:

- To create a bootable USB flash drive by using Etcher:
 1. Download the latest image of the Test Agent Appliance from the [Software Download](#) page.
 2. Download and install [Etcher](#) on your device to create a bootable USB flash drive.
 3. Insert a USB flash drive with at least 4.1 GB of free space into your hardware.
 4. Launch Etcher, and select the Test Agent Appliance image from the hard disk of your x86 hardware.
 5. Select the USB flash drive where you want to write the image.
 6. Click **Flash** to start the process.

Once the image has been flashed, you can eject your USB flash drive.

7. Eject the USB flash drive.

The USB flash drive is now ready to be used as a bootable disk.

For more information on Etcher, see [Etcher documentation](#).

- To create a bootable USB flash drive by using **Win32 Disk Image**, see [Create a Bootable USB flash drive Using a Windows Device](#).
- To create a bootable USB flash drive by using **dd** command on Max or Linux, see [Create a Bootable USB flash drive Using a MAC OS X](#).



WARNING: Make sure that the USB flash drive has a memory of 4.1 GiB for temporary installation files, and any existing content on the USB flash drive is erased.

Boot the custom x86 hardware from USB flash drive

You can boot the custom x86 hardware from USB flash drive by using the BIOS boot manager.

To boot the custom x86 hardware from USB flash drive:

1. Insert the bootable USB flash drive into a USB port on your custom x86 hardware.
2. Power on the x86 hardware, and press DEL, ESC, F2, or F12 depending on the hardware to access the BIOS menu.

You can now access the BIOS boot manager to check the USB selection.

3. Go to **Boot Order** settings and move the USB flash drive to the top of the priority list.
4. Select **Save** to update the settings and reboot.

The boot process takes up to 20 seconds, and once complete a login prompt is displayed.

5. Log in to the system by using the system credentials. The username is 'admin' and the password is 'admin'.
6. Navigate to **Utilities** and select **Install to disk**.

Once the installation is complete, you are prompted to remove the USB flash drive and reboot the hardware.



NOTE: Remove the USB flash drive before you reboot your x86 hardware.

7. Reboot your x86 hardware.

You have successfully installed Test Agent Appliance on a x86 hardware, and the custom x86 hardware now boot up as a Test Agent Appliance.

Before you can use the Test Agent Appliance, you must register Test Agent Appliance with the Routing Director.

Register a Test Agent Appliance from the local console

When you reboot the custom x86 hardware, the Test Agent Admin Menu appears. To use the Test Agent Appliance, you must first configure the management interface to enable connectivity with Routing Director. Once you configure, you can proceed with registration.

Configure the management interface on Test Agent Appliance

To configure the management interface on a Test Agent Appliance:

1. Reboot the custom x86 hardware to access the Local Console.

The text-based Test Agent Admin Menu appears automatically when you reboot the custom x86 hardware.

2. Select **Configure management**, and click **OK**.

A Select Management Interface page appears.

3. Select the interface for management traffic, and click **OK**.

The Select Interface speed page appears.

4. Select the speed and duplex settings for the selected management interface, and click **OK**.

The page appears asking whether you want to use a VLAN tag for the management interface.

5. (Optional) Specify the VLAN tag. If your network setup requires VLAN tagging, select **Yes** and specify the appropriate VLAN ID.

6. Select one of the following IPv4 address types, and click **OK**.

- None—To not configure an IPv4 address type.
- Static—To manually specify the IPv4 address.
- DHCP—To automatically assign IPv4 address.

7. Select one of the following IPv6 address types, and click **OK**.

- None—To not configure an IPv6 address type.
- SLAAC—To automatically configure a stateless address.

- Static—To manually specify the IPv6 address.
- DHCP—To automatically assign IPv6 address.

Once you configure an IPv6 address, a page is displayed asking whether you want to enable IPv6 support for the NTP client.

8. Specify NTP servers as comma-separated list for accurate time synchronization, and click **OK**.

The page appears asking whether the information you specified is correct.

9. Verify the information, and click **Yes**.

Now, you are taken back to the Test Agent Admin Menu where you can register Test Agent Appliance with Routing Director.

Register a Test Agent Appliance by using Secret

To register a Test Agent Appliance by using Secret:

On the Routing Director GUI—

1. Navigate to the Test Agents (**Inventory > Active Assurance > Test Agents**) page.
2. Click the **Add Test Agent** button.

The Add Test Agent page appears.

3. Select Test Agent Appliance as the type of Test Agent.
4. Specify the Name, Tags, and Description for the Test Agent.
5. Click **Create**.

A message confirms that Routing Director has successfully added the Test Agent Appliance, and you are redirected to the Test Agent Details (**Inventory > Active Assurance > Test Agents > Test Agent Details**) page.

6. On the **Download & Run** tab, copy the secret key generated by Routing Director by using the **Copy** button next to the Secret text box. You can use this secret key to establish a connection with Routing Director.

On the Test Agent Admin Menu—

1. Configure the management interface as described in ["Configure the management interface on Test Agent Appliance" on page 937](#).
2. Select **Register/Connect** from the text-based Test Agent Admin Menu.

The Register/Connect Menu page appears.

3. Select **Register using secret**, and click **OK**.

The Enter registration information page appears.

4. Specify the following details, and click **OK**.

- IP address of the Routing Director Test Agent Gateway.
- TCP port number that connects to Routing Director. The default TCP port number is 6800.
- Secret that you copied from the **Download & Run** tab.

A confirmation message appears displaying that the registration is successful.

Once the registration is successful, Test Agent Appliance is listed on the Test Agents (**Inventory > Active Assurance > Test Agents**) page on the Routing Director GUI.



NOTE: If the Routing Director uses a self-signed certificate, navigate to **Utilities > SSL Certificate Validation** on the Test Agent Admin Menu and disable the SSL certificate validation.

Register a Test Agent Appliance by using username and password

You can register a Test Agent Appliance by using username and password on the text-based Test Agent Admin Menu.



NOTE: The preferred method for registering a Test Agent Appliance is by using a secret.



WARNING: The username and password authentication does not work if Single Sign-On (SSO) and Two-Factor Authentication (2FA) are enabled.

To register a Test Agent Appliance by using username and password:

1. Select **Register/Connect** from the Test Agent Admin Menu.

The Register/Connect Menu page appears.

2. Select **Register using user/password** and click **OK**.

The Enter registration information page appears.

3. Specify the following details, and click **OK**.

- IP address of the Routing Director Test Agent Gateway.
- TCP port number that connects to Routing Director. The default TCP port number is 6800.
- Email address associated with your Routing Director account.
- Name for Test Agent Appliance.

4. Enter your Routing Director account password for authentication, and click **OK**.

A Select organization page appears.

5. Select your organization from the list, and click **OK**.

A confirmation message appears displaying that the registration is successful.

Once the registration is successful, Test Agent Appliance is listed on the Test Agents (**Inventory > Active Assurance > Test Agents**) page on the Routing Director GUI.

Install Test Agent Appliance on Proxmox

Proxmox is an open-source virtualization platform where you can deploy a Test Agent Appliance as a virtual machine (VM). This method helps you to run a Test Agent Appliance in a virtualized environment instead of dedicated physical hardware,

To install Test Agent Appliance on Proxmox server:

1. Create a VM without an operating system.

a. Log in to the Proxmox Web UI.

b. Click the **Create VM** button from the left-nav bar.

The Create: Virtual Machine page appears.

c. In the General tab, specify the VM ID to identify the VM.

d. In the OS tab, select **Do Not Use Any Media** to prevent the VM from booting from a CD-ROM or ISO file.

e. In the Disk tab, click the **Delete** (trash can) icon to delete the default disk.

f. In the CPU, Memory, and Network tabs, specify details based on the VM requirements, and click **Finish**. For more information, see [Proxmox VM settings](#).



NOTE: For a Test Agent, the requirements are:

- CPU— Minimum 1 vCPU (Recommended—2 vCPU)
- Memory—Minimum 256 MB RAM (Recommended—4 GB RAM)
- Disk Space—4.1 GB Disk

2. Create a directory to store qcow2 images.

Before you begin, download the latest image of the Test Agent qcow2 from the [Software Download](#) page.

- In the Proxmox UI, click **Host** from the left-nav bar and click Shell to open the terminal.
- Use the following command to create a directory:

```
sudo mkdir /var/lib/vz/template/qcow
```

- Navigate to created directory and run the following command to download the image:

```
wget ta-ng.qcow2
```

After you download, proceed to import the qcow2 image.

3. Import the qcow2 image.

On the Proxmox Shell, run the following command:

```
qm importdisk <id> <paa-test-agent_5.0.0.4.qcow2 local-lvm>
```



NOTE: The <values> used in the commands are placeholders. You must replace them with required values based on your configuration need.

4. Attach the qcow2 virtual disk to the VM.

- Select the VM you created and navigate to the **Hardware > Hard disk**.
- Select the newly imported QCOW2 disk that is listed under Unused Disk and click the **Edit** button.

The Add: Unused Disk page appears.

- Select the bus type **VirtIO Block** from the drop-down list and click **Add**.

You have successfully attached qcow2 disk to the VM as a VirtIO Block device.

5. Change the boot order.

- a. Navigate to the **Options > Boot Order**.
- b. Select Boot Order and move the VirtIO Block Disk to the top.
- c. Click **OK**.

6. Start the VM.

- a. Navigate to the **Console** menu.
- b. Click **Start Now** to boot the Test Agent Appliance.

You have successfully deploy the Test Agent Appliance on Proxmox. You can use the cloud-config metadata to establish a communication with Routing Director.

Format of cloud-config metadata

On virtualization platforms, cloud-config metadata (user data) can be specified. During deployment of a virtual machine, cloud-init uses cloud-config (a YAML-based file format) to automate the initialization of virtual machines. You can customize cloud-config metadata to define network configurations and application-specific parameters. When you deploy a Test Agent Appliance on Amazon Web Services or OpenStack, the cloud-config metadata helps to automate Test Agent Appliance configuration. It establishes secure communication with the management server registering Test Agent Appliance with Routing Director.

You can use the following cloud-config file as reference:

```
#cloud-config
paa_test_agent:

  # Registration using secret
  server: "<server_address>"
  port: 6800
  secret: "<registration_secret>"

  # Disable/enable SSL certificate validation
  ssl_no_check: false # (Optional) Set to true to disable SSL checks (not recommended)

  # Optional HTTP proxy settings
```

```

http_proxy:
  address: "<proxy_address>" # (Required if using a proxy) Proxy server address
  port: <proxy_port> # (Required) Proxy server port
  username: "<proxy_username>" # (Optional) Proxy authentication username
  password: "<proxy_password>" # (Optional) Proxy authentication password

# User account passwords
user_passwords:
  admin: "<admin_password>" # (Optional) Password for the admin user
  root: "<root_password>" # (Optional) Password for the root user

# Network management configuration
management:
  interface: "eth0" # (Optional) Network interface (default: eth0)
  mtu: <mtu_value> # (Optional) MTU size (default: None)
  address_type: "dhcp" # (Optional) IP assignment type: 'none', 'dhcp', or 'static'
  (default: dhcp)
  ip: "<static_ip>" # (Required if using 'static' address type) Static IPv4 address
  dns: ["<dns_server_1>", "<dns_server_2>"] # (Optional) List of DNS servers
  default_gateway: "<gateway_ip>" # (Optional) Default IPv4 gateway
  ntp: ["time.google.com"] # (Optional) NTP server(s) (default: time.google.com)
  address6_type: "dhcp" # (Optional) IPv6 assignment type: 'none', 'dhcp', 'static', or
  'slaac' (default: dhcp)
  ip6: "<static_ipv6>" # (Required if using 'static' address type) Static IPv6 address
  dns6: ["<dns6_server_1>", "<dns6_server_2>"] # (Optional) List of IPv6 DNS servers
  default_gateway6: "<gateway_ipv6>" # (Optional) Default IPv6 gateway
  ntp_allow_ipv6: false # (Optional) Enable IPv6 NTP (default: false)

# Test Agent application settings
ta_application:
  log_level: "warn" # (Optional) Logging level: 'debug', 'info', 'warn', 'error',
  'critical' (default: warn)

```



NOTE: The <values> used in the commands are placeholders. You must replace them with required values based on your configuration need.

RELATED DOCUMENTATION

Active Assurance Workflow | 891

Tests and Monitors

IN THIS CHAPTER

- [Tests and Monitors Overview | 944](#)
- [About the Measurement Designer Page | 955](#)
- [Create a Test | 958](#)
- [Create Input Variables | 1097](#)
- [About the Tests Page | 1100](#)
- [About the *Test-Name* Page | 1104](#)
- [Create a Monitor | 1110](#)
- [About the Monitors Page | 1247](#)
- [About the *Monitor-Name* Page | 1253](#)
- [View Stream Details | 1258](#)

Tests and Monitors Overview

IN THIS SECTION

- [Measuring Metrics Using Active Assurance Tests | 945](#)
- [Measuring Metrics Using Active Assurance Monitors | 946](#)
- [Supported Protocols | 946](#)
- [Benefits of Using Active Assurance Tests and Monitors | 954](#)

Monitoring key performance indicators (KPIs) such as the response time, congestion, and reachability is important to gauge the quality of your network. To evaluate the KPIs, Active Assurance enables you to create and run *Test* and *Monitors* to measure common metrics such as delay, delay variance (jitter), packet loss, round-trip response time, and so on.

A Test is a set of verifications that is performed by one or more Test Agents for a finite amount of time. For more information on Tests, see ["Measuring Metrics Using Active Assurance Tests" on page 945](#).

A Monitor is a set of verifications that is performed by one or more Test Agents for an infinite amount of time. For more information on Monitors, see ["Measuring Metrics Using Active Assurance Monitors" on page 946](#).

Both Tests and Monitors entails *Tasks*. The Task contains the configuration to measure a specific metrics by using certain protocols. For information on the protocols that are supported in Routing Director, see ["Supported Protocols" on page 946](#).

When you create and run a Test or Monitor, the Test Agent receives the measurement configuration that you have defined in the Test or Monitor. Based on the protocols that you have defined during the Task creation, the Test Agent downloads the required protocols from Routing Director, starts collecting the measurements, and shares the data in the form of streams that you can view in the Routing Director GUI. This information helps you to gain insights into your network's health, performance, and quality.


Measuring Metrics Using Active Assurance Tests

A Test is a set of verifications that are performed by one or more Test Agents for a finite amount of time. A Test contains one or more *Steps* and *Tasks*, which validates whether KPIs that are defined in the Tests can be considered as operational (pass) or not (fail).

A Test can consist of one or more Steps that are executed sequentially. Each Step can consist of one or more Tasks that run concurrently. A Task contain the configuration to measure specific metrics.

[Figure 48 on page 945](#) illustrates the relation between a Test, Steps, and Tasks in Routing Director.

Figure 48: Components in an Active Assurance Test



Test measures metrics in a time-bound manner and therefore you need to set the duration for each Step which determines the duration of the entire Test.

A Test delivers a binary result—Pass or Fail.



NOTE: If any one Task with the Test fails, then Test status is displayed as Fail.

Measuring Metrics Using Active Assurance Monitors

A Monitor is a set of verifications that are performed by one or more Test Agents for an infinite amount of time. A Monitor can contain multiple Tasks, which run in parallel and continuously monitors the KPIs that are defined in the Monitor. A Task contains the configuration to measure specific metrics.

[Figure 49 on page 946](#) illustrates the relation between a Monitor and Tasks in Routing Director.

Figure 49: Components in an Active Assurance Monitor



Monitors measure the metrics indefinitely until you decide to stop the Monitor. There is no limited time duration for which you want to stop a Monitor. Since you can start and stop a Monitor at will you need not set the duration.

A Monitor delivers a time-framed result as it runs continuously.

Supported Protocols

Routing Director enables you to configure the protocols listed in [Table 263 on page 947](#) for a Task.

Table 263: List of Protocol Tasks

Protocol	Description
DNS	The DNS task is used to test and monitor DNS servers. When a DNS task starts, the test agents send a request to resolve a lookup address, and collect statistics on response times. DNS primarily uses User Datagram Protocol (UDP) on port number 53 to serve requests. DNS queries consist of a single UDP request from the client followed by a single UDP reply from the server. Running a DNS task provides information about the response times of your DNS servers from different locations. High DNS response times translate into high response times for all services that use DNS to resolve IP addresses, such as web surfing. This task is applicable for both IPv4 and IPv6.
HTTP	The HTTP Task is used to test or monitor HTTP servers. Running an HTTP Task checks the performance of a website or web application of the web server and of the network between the web server and the Test Agent. You can request web pages and verify response codes from distributed locations inside or outside of your network. When an HTTP Task starts, the Test Agents make an HTTP Get request towards the specified URL and fetch the response. The task does not render HTML pages. Hence, test agents do not make additional requests for linked resources, (images, CSS files, and so forth). Measured parameters include TCP connect time, time until first byte received, time until last byte received, and download speed. This task is applicable for both IPv4 and IPv6.

Table 263: List of Protocol Tasks *(Continued)*

Protocol	Description
IPTV MPEG	The IPTV MPEG Task is used to monitor the IPTV channel quality. Running an IPTV MPEG Task provides information about the collecting MPEG loss, PCR jitter, rate, packet loss, and continuity count errors at Test Agent end-points. When the Task starts, the Test Agents join the channels by sending IGMP join messages. On receiving the MPEG streams, the Test Agents continuously measures quality.
Netflix Speedtest	The Netflix Speedtest Task is used to assess the performance of a Netflix speed test transaction. Running a speed test provides information on the bandwidth, latency, or threshold violations, if any. Routing Director instructs Test Agents to download Netflix test segments over HTTPS, so that your network can handle Netflix streaming optimally.
OTT-HLS	The OTT-HLS Task is used to integrate OTT service monitoring and selecting the highest possible quality to avoid buffering. When running an OTT-HLS Task, Test Agent parses the manifest file and starts downloading the video segments. The algorithm adapts to the current network conditions and selects the highest bit rate to avoid buffering.

Table 263: List of Protocol Tasks *(Continued)*

Protocol	Description
Ping	The Ping task initiates a task that checks connectivity to the remote device. Running a ping task provides information about the delay, delay variance (jitter), packet loss of the connection to the remote host. The Ping tool uses Internet Control Message Protocol (ICMP) or UDP to initiate a single request from the test agent to the host, followed by a single response from the host. This task works with both IPv4 and IPv6.
RPM CISCO TWAMP Reflector	The RPM Cisco TWAMP Reflector task measures network performance by running TWAMP (Two-Way Active measurement protocol) tests specifically for Cisco devices. This test evaluates metrics like round-trip delay, jitter, and packet loss between Test Agents (senders) and Cisco devices (responders). NOTE: We recommend that you configure RPM Cisco TWAMP Reflector using API as configuring RPM Cisco TWAMP Reflector using the Routing Director GUI is a Beta feature.
RPM HTTP	The RPM HTTP Task checks connectivity to the remote device. Use this Task if you are using ACX, PTX or MX device.

Table 263: List of Protocol Tasks *(Continued)*

Protocol	Description
RPM ICMP	Use this Task if you are using ACX, PTX or MX device. You can measure one-way measurements for ICMP timestamp probes that includes information on • Minimum, maximum, standard deviation, and jitter measurements for egress and ingress times • Number of probe responses received • Percentage of lost probes • Number of probes sent
RPM TCP	Use this Task if you are using ACX, PTX or MX device.
RPM TWAMP	Use this Task if you are using ACX, PTX or MX device. The two way active measurement protocol (TWAMP) facilitates the measurement of two-way or round-trip network performance metrics. Running a TWAMP Task provides information about the round-trip delay, delay variance (jitter), and packet loss. The Session-Initiator creates TWAMP test packets and sends to the Session-Reflector in the TWAMP server, and the Session-Reflector sends back a measurement packet when a test packet is received. TWAMP uses TWAMP-Control protocol to perform a handshake between initiator and reflector.

Table 263: List of Protocol Tasks *(Continued)*

Protocol	Description
RPM UDP	Use this Task if you are using ACX, PTX or MX device. This Test checks if your network is good enough for quality-demanding services such as client-server applications and video conferencing. When a UDP Task starts, the Test Agents will generate traffic at the rate you specify. The rate is the Layer 2 Ethernet rate, also known as the Committed Information Rate (CIR). It includes the Ethernet headers with the CRC checksum but not the Frame Gap, Preamble, or Start of Frame Delimiter. The UDP flow sent by the sender Test Agent includes timestamps and sequence numbers, so that the receiving Test Agent can calculate one-way delay, jitter, and packet loss.
TCP	The TCP Task is used to assess the network performance for client-server applications by sending TCP sessions between Test Agents. Running a TCP Task provides information about the latency, jitter, throughput bandwidth, and monitor network congestion. When a TCP Task starts, the client Test Agents send TCP packets to server Test Agents, and receives measurement information from server Test Agents. This gives an overall view of the network.

Table 263: List of Protocol Tasks *(Continued)*

Protocol	Description
TWAMP Reflector	The TWAMP Reflector is a component of TWAMP. It receives test packets from the Session-Initiator and reflects them back. Running a TWAMP Reflector helps in evaluation of the quality of the network and provides information on how data is transmitted in two-directions; between Session-Initiator and Session-Reflector. This bidirectional communication enables the measurement of performance metrics such as round-trip delay, delay variance (jitter), and packet loss. This Task is applicable for both IPv4 and IPv6.
TWAMP/TWAMP Light	• TWAMP—The two way Active measurement protocol facilitates the measurement of two-way or round-trip network performance metrics. Running a TWAMP Task provides information about the round-trip delay, delay variance (jitter), and packet loss. The Session-Initiator creates TWAMP test packets and sends to the Session-Reflector in the TWAMP server, and the Session-Reflector sends back a measurement packet when a test packet is received. TWAMP uses TWAMP-Control protocol over TCP to perform a handshake between initiator and reflector. • TWAMP Light—The two way Active measurement protocol light is a simplified version of TWAMP. It is a stateless version of TWAMP where test parameters are predefined instead of negotiated. All test packets received by the server on a test port are reflected back and forgotten right away. Running a TWAMP Light Task provides information about the round-trip delay, delay variance (jitter), and packet loss. Unlike TWAMP, TWAMP Light does not use TWAMP-Control protocol. This Task is applicable for both IPv4 and IPv6.

Table 263: List of Protocol Tasks *(Continued)*

Protocol	Description
UDP	The UDP Task is used to evaluate the network quality for client-server applications and video conferencing. Running a UDP Task measures one-way delay, jitter, packet loss, and misorders of packets. When a UDP task starts, the client Test Agents send packets that include timestamps and sequence numbers to the server Test Agents at the rate you specify. In this way, the server Test Agent can calculate one-way delay, jitter, and packet loss.
Y.1731 DM	Y.1731 Delay Measurement (DM) task is used to evaluate the network performance with respect to time. It is a service OAM (Operations, Administration, and Maintenance) mechanism defined by the ITU-T to monitor network performance. Running a Y.1731 DM task measures delay and delay variance of packets. The local maintenance end point (MEP) sends packets that has timestamps to the remote maintenance end point (MEP) and vice-versa. The Test Agents then calculate one-way delay as well as two-way delay measurements. For more information on Y.1731 Ethernet Service OAM, see ITU-T Y.1731 Ethernet Service OAM Overview.

Table 263: List of Protocol Tasks *(Continued)*

Protocol	Description
Y.1731 SLM	Y.1731 Synthetic Loss Measurement (SLM) task is used to measure the frame loss across the network. It is a service OAM (Operations, Administration, and Maintenance) mechanism defined by the ITU-T to monitor network performance. Running a Y.1731 Synthetic Loss Measurement (SLM) task measures the number of frames lost during transmission. The local MEP sends synthetic frames to the remote MEP and compares the number of frames transmitted with the number received. The ratio of lost frames to the total number of frames that are sent is calculated to measure the frame loss on the network. For more information on Y.1731 Ethernet Service OAM, see ITU-T Y.1731 Ethernet Service OAM Overview.

Benefits of Using Active Assurance Tests and Monitors

- Enables you to analyze reachability, congestion, and response time in networks to support business critical applications.
- Access historical view of errors in your network.
- Customize parameters and thresholds to reflect the optimum network performance that you require.

RELATED DOCUMENTATION

[About the Tests Page | 1100](#)

[About the Monitors Page | 1247](#)

About the Measurement Designer Page

IN THIS SECTION

- [Tasks You Can Perform](#) | 955

You can access the Measurement Designer page in one of the following ways:

- Click **Observability** > **Active Assurance** > **Measurement Designer**.
- Click **Observability** > **Active Assurance** > **Tests** > **Add (+)**.
- Click **Observability** > **Active Assurance** > **Monitors** > **Add (+)**.

You (superusers or network administrators) can use the Measurement Designer page to create a Test or a Monitor, configure Steps and Tasks, edit a Test or a Monitor name, edit a Step name and its duration, switch between Test and Monitor modes, remove a Step or a Task, and so on. You can also zoom in, zoom out, fit Steps to page size, and lock or unlock the Steps that you are configuring.

Tests and Monitors are used to measure common metrics such delay, delay variance (jitter), packet loss, round-trip response time, and these collect metrics by using protocol configurations called Tasks. For more information on Tests and Monitors, see ["Tests and Monitors Overview" on page 944](#).

Tasks You Can Perform

You can perform the following tasks on the Measurement Designer page:

- Create a Test. See ["Create a Test" on page 958](#).
- Create a Monitor. See ["Create a Monitor" on page 1110](#).
- Specify a name for a Test or a Monitor—You need to add a name for the Test or the Monitor that you are creating.

If you are creating a Test, click the **Sequence Name** or click the **Edit** (pencil) icon next to **Sequence Name** and enter a name.

If you are creating a Monitor, click the **Monitor Name** or click the **Edit** (pencil) icon next to **Monitor Name** and enter a name.

Sequence Name is the placeholder name for a Test and **Monitor Name** is the placeholder name for a Monitor.

- Add a Step—A Step is a collection of protocol Tasks.

In Test creation mode, you can add one or more Steps by clicking **+ Add Step**.

In Monitor creation mode, there can be only one Step, and it is available by default.

- Add a Task—Each Step can include one or more Tasks. You can add multiple Tasks from the Task selector page displayed by default, or by clicking **+ Add Task**. On the Task page, you can configure different parameters and set evaluation criteria.
- Switch between Test and Monitor—You can switch from the Test to the Monitor creation mode by using the **Monitor** button and vice versa. If you want to switch from the Test creation mode to Monitor creation mode, ensure that you have configured only one Step as a Monitor can only have a single Step with concurrent Tasks. If you have configured more than one Step, the Steps other than the first one will not be moved to the Monitor creation mode. If you want to create a Monitor based on a Step that is not the first Step, you can move the desired Step to the first position and then switch to the Monitor mode.

The name you configured for Test and Monitor will remain the same when you switch between modes.

The duration you configured for a Step in Test will be deleted when you switch to the Monitor mode as a Monitor runs indefinitely.

- Manage input variables for Tests—You can manage the input variables you created by clicking the **Global** icon on the right side of the page. For more information on input variables, see ["Create Input Variables" on page 1097](#).
- Configure Test and Monitor settings—You can configure additional settings for a Test and a Monitor. Click the **Test Settings** (gear) icon or the **Monitor Settings** (gear) icon to,
 - Specify the wait duration (in seconds)—You can specify the maximum wait duration. After the Test Agent completes a Step, Routing Director waits for the Test Agents to configure or download the plugins or for the Test Agents to be ready to run the Steps.

By default, the Step wait duration is 60 seconds. The Step wait duration has no effect on the Test duration.

You cannot add maximum wait duration for a Monitor as the Monitor can only contain one Step.

- (Optional) Specify a short description a the Test or a Monitor.
- Add **Tags**—Specify a key-value pair in the key:value format.

You can configure Tags for a Test or a Monitor in the key:value format to provide additional information about the Test or Monitor you are configuring. A tag is a key-value pair in which the key signifies a category for which you configure a value. The value is an identifier for the category.

Examples of key-value pairs are device and device name (edgedevice:acx7000), site and site name (site:bangalore).

- **Save a Test as a template**—You can save a Test as a template and reuse this template for creating Test instances in a consistent and reproducible manner. Click the **Save as Template** button on the top-right corner of the page.

A Save as Test Template page appears asking you to confirm the save operation.

A message is displayed confirming that the Test template is saved successfully, and the Test template you created is listed on the Test Templates (**Inventory > Active Assurance > Test Templates**) page. For more information, see ["About the Test Templates Page" on page 1288](#).

- **Run a Test or a Monitor**—You can run a Test or a Monitor. After you configure a Test and a Monitor, click the **Run** button.



NOTE: If a Step does not contain any Task, the **Run** button will be disabled.

A successful message is displayed confirming that you have successfully run a Test or a Monitor, corresponding measurements are generated and specified configurations are pushed to the Test Agents. Now, you are redirected to the *Test-Name* (**Observability > Active Assurance > Tests > Test-Name**) page or the *Monitor-Name* (**Observability > Active Assurance > Monitors > Monitor-Name**) page where you can view the Tests or Monitors you run.

- Additionally, you can perform the following tasks from the bottom-left corner of the page:
 - **Zoom in**—You can click the **Zoom in** (plus) icon to zoom in the Steps that you are configuring.
 - **Zoom out**—You click the **Zoom out** (minus) icon to zoom out the Steps that you are configuring.
 - **Maximize**—You can click the **Fit view** (box) icon to fit Steps to the page size.
 - **Toggle interactivity**—You can click the **Toggle interactivity** (lock) icon to lock or unlock the Steps.

When you lock the Toggle interactivity icon, you can drag the Steps anywhere on the Measurement Designer page but you cannot edit the Step.

When you unlock the Toggle interactivity icon, you cannot drag the Step anywhere on the Measurement Designer page but you can edit the Step.

Create a Test

IN THIS SECTION

- [Configure Test settings | 958](#)
- [Configure Task Parameters | 961](#)
- [Native Plug-ins | 963](#)
- [RPM Plug-ins | 1048](#)
- [Configure Evaluation Criteria | 1095](#)

Configure Test settings

Tests measure metrics in a time-bound manner to produce a static measurement for the given duration. Tests have two sub components namely, Steps and Tasks. A Test is a collection of Steps, and a Step is a collection of Tasks. In Tests, Tasks under Steps are run parallel and each Step runs for the duration you configure. For more information, see "[Tests and Monitors Overview](#)" on page 944.

To create a Test:

1. Do one of the following to access the Measurement Designer page:
 - a. Access Measurement Designer page through the Tests page.
 - i. Navigate to the Tests (**Observability** > **Active Assurance** > **Tests**) page.
The Tests page appears.
 - ii. Click the **Add (+)** icon.
The Create new page appears.
The Create new page displays **All**, **Tests**, and **Monitors** mode, and the Tests mode is enabled by default.
 - iii. Select the **+ Create blank Test**.
The Measurement Designer page appears.
 - b. Access Measurement Designer page directly.
 - i. Navigate to the Measurement Designer (**Observability** > **Active Assurance** > **Measurement Designer**) page.

The Create new page appears.

The Create new page displays **All**, **Tests**, and **Monitors** modes, and the **All** mode is enabled by default.

- ii. Select the **Tests** mode.
- iii. Select the **+ Create blank Test**.

The Measurement Designer page appears.

2. Specify a name for the Test.

Click the **Sequence Name** or click the **Edit** (pencil) icon to enter a name for the Test in the **Sequence Name** text box.

Sequence Name is the placeholder name for the Test.

3. Click **+ Add Step** to add Steps for a Test.



NOTE: In the Measurement Designer page, the first Step and the Task selector are available by default in the GUI.

For the Step, you can do the following:

- (Optional) Edit the name of a Step—Click the **Edit** (pencil) icon near the *Step Name* and specify a name for the Step.
- (Optional) Edit the duration of a Step—Click the **Edit** (pencil) icon near the **Duration** (clock) icon to edit the duration.

The default duration for a Step is set to 30 seconds. The duration range is 30s through 86400s.

- (Optional) Delete a Step—To delete a step, do one of the following:
 - Click the horizontal ellipsis (three dots) in the Step and click **Remove**.
 - Drag the Step anywhere on the screen.

A confirmation message appears asking to confirm if you want to remove the Step. Click **Yes**.



NOTE: Fields marked with an asterisk (*) are mandatory.

4. Select a Task from the Task selector page displayed by default, or click **+ Add Task** to add Tasks for a Step. The Tasks page displays the list of protocols that can be configured for a Test.

In the Tasks page, do one of the following:

- Select a Task from the Tasks page. For example, click DNS to add DNS as one of the Tasks in the Step.
- Click the **Add** (+) icon next to the Task.
- Drag and drop the Task inside the Step that you are configuring.

Select the Close (X) icon to close the Tasks page.

For a Task, you can do the following:

- (Optional) Edit the name of a Task—To edit a Task, do one of the following:
 - Click the **Edit** (pencil) icon to specify a name for the Task in the *Task-Name* text box.
 - Click the *Task-Name* and specify a name for the Task in the *Task-Name* text box.

By default, the plugin name is displayed. If you do not edit, the default name will be used.

- (Optional) Delete the Task—To delete a Task, do one of the following:
 - Click the down arrow in the Task box and click **Remove**.
 - Drag the Task anywhere on the empty portion of the screen.

A confirmation message appears asking to confirm if you want to remove the Task. Click **Yes**.

- Configure the parameters for a Task—Click the **Settings** (gear) icon on the Task box to configure a Task. The **Step Name** page appears and displays the Task you added. This page includes the following tabs:
 - Parameters tab—Configure parameters for the Tasks that you have added. For more information on parameters that you can configure for a Task, see [Table 264 on page 961](#).

You can also configure input variables for each parameter so that the same values are consistently used across different Steps. For more information on input variables, see ["Create Input Variables" on page 1097](#).
 - Evaluation criteria tab—The Evaluation Criteria for each Task is added by default. You can configure customized evaluation criteria for metrics by using threshold expressions. On violation of any configured expression, Routing Director generates events. For more information on metrics that you can configure, see [Table 283 on page 1095](#).

5. Click + Add Task to add Task for a Step.

A Test can include one or more Steps, and a Step can include one or more Tasks. Based on your requirements, repeat ["Step 3" on page 959](#) and ["Step 4" on page 960](#).

6. Click the **Global** icon on the right side of the Measurement Designers page to manage input variables. For more information on input variables, see ["Create Input Variables" on page 1097](#).

7. Click **Test Settings** (gear) icon on the right side of the page to,

- Specify the wait duration (in seconds)—You can specify the maximum wait duration. After the Test Agent completes a Step, Routing Director waits for the Test Agents to configure or download the plugins or for the Test Agents to be ready to run the Steps.

By default, the Step wait duration is 60 seconds. The Max wait duration has no effect on the Test duration.

- (Optional) Specify a short description for the Test.
- (Optional) Add **Tags**—Specify a key-value pair in the key:value format.

You can configure Tags for the Test in the key:value format to provide additional information about the Test you are configuring. A tag is a key-value pair in which the key signifies a category for which you configure a value. The value is an identifier for the category. Examples of key-value pairs are device and device name (edgedevice:acx7000), site and site name (site:bangalore).

8. Click **Run**.

A message confirming that the Test is created successfully appears and you are redirected to the *Test-Name* (**Observability > Active Assurance > Tests > Test-Name**) page. On the *Test-Name* page, Test status is displayed as **Running** indicating that the Test is in progress. The status changes to **Completed** when the Test is complete and you can view the details of the Test. See ["About the Test-Name page" on page 1104](#) for more information about the Test details.



NOTE: The **Run** button is disabled until you add a Task under a Step. Also, You cannot run a Test if you have not specified the **Test Name**.

Configure Task Parameters

Table 264: Tasks and its parameters

Tasks	Description
DNS	For information on the parameters that you can configure for DNS, see Table 265 on page 963 .
HTTP	For information on the parameters that you can configure for HTTP, see Table 266 on page 971 .

Table 264: Tasks and its parameters *(Continued)*

Tasks	Description
Ping	For information on the parameters that you can configure for ping, see Table 267 on page 977 .
RPM CISCO TWAMP Reflector	For information on the parameters that you can configure for ping, see Table 268 on page 985 . NOTE: We recommend that you configure RPM Cisco TWAMP Reflector using API as configuring RPM Cisco TWAMP Reflector using the Routing Director GUI is a Beta feature.
IPTV MPEG	For information on the parameters that you can configure for IPTV MPEG, see Table 269 on page 991 .
Netflix Speedtest	For information on the parameter that you can configure for Netflix Speedtest, see Table 270 on page 997 .
OTT - HLS	For information on the parameters that you can configure for OTT-HLS, see Table 271 on page 1002 .
TCP	For information on the parameters that you can configure for TCP, see Table 272 on page 1009 .
UDP	For information on the parameters that you can configure for UDP, see Table 273 on page 1019 .
TWAMP/TWAMP Light	For information on the parameters that you can configure for TWAMP/TWAMP Light, see Table 274 on page 1031 .
TWAMP Reflector	For information on the parameters that you can configure for TWAMP Reflector, see Table 275 on page 1044 .

Table 264: Tasks and its parameters *(Continued)*

Tasks	Description
RPM HTTP	For information on the parameters that you can configure for RPM HTTP, see Table 276 on page 1048 .
RPM PING	For information on the parameters that you can configure for RPM PING, see Table 277 on page 1054 .
RPM TCP	For information on the parameters that you can configure for RPM TCP, see Table 278 on page 1061 .
RPM TWAMP	For information on the parameters that you can configure for RPM TWAMP, see Table 279 on page 1067 .
RPM UDP	For information on the parameters that you can configure for RPM UDP, see Table 280 on page 1075 .
Y.1731 DM	For information on the parameters that you can configure for Y.1731 DM, see Table 281 on page 1082 .
Y.1731 SLM	For information on the parameters that you can configure for Y.1731 SLM, see Table 282 on page 1088 .

Native Plug-ins

Table 265: DNS Parameters

Parameter	Description
General	

Table 265: DNS Parameters (Continued)

Parameter	Description
Clients*	Select one or more Test Agent interfaces on which you want to run the Test. To select one or more Test Agents: 1. Click the Clients text box. The Select interfaces page appears. 2. Select one or more Test Agents. For information on filtering Test Agents, see "Search for Test Agent interfaces by using Filters" on page 964. 3. Click OK. NOTE: On the Select Test Agent/Interfaces page, you can filter the Test Agent interfaces based on the following criteria: • Interface name or the Test Agent name, or both. Based on your requirement, you can select or deselect Test Agents or Interfaces. • IPv4 address (with or without IPv4 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv4 or No IPv4 to view or hide the interfaces that are associated with the IPv4 family and the interfaces that are not associated with the IPv4 family. • IPv6 address (with or without IPv6 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv6 or No IPv6 to view or hide the interfaces that are associated with the IPv6 family and the interfaces that are not associated with the IPv6 family. • Deployment type of a Test Agent:

Table 265: DNS Parameters *(Continued)*

Parameter	Description
	• Application—Filter out the Test Agents deployed in a device as an application. • Device—Filter out the Test Agents associated with a device. Based on your requirement, you can select or deselect Application or Device to view or hide the Test Agents deployed in a device as an application or the Test Agents associated with a device. • Connection types between the interface and the router: • Link—Indicates that the status is Up and there is an active connection for the interface. • No Link—Indicates that the status is Down and there is no active connection for the interface. Based on your requirement, you can select or deselect Link or No Link. • Connection statuses of the Test Agent: • Online—The Test Agent is connected to Routing Director. • Offline—The Test Agent is not connected to Routing Director. When the interface supports RPM plugins, the interface contains <i>(On device)</i> at the end of its name indicating the Device Online status. When the interface supports only the normal plugins, the interface will not display <i>(On</i>

Table 265: DNS Parameters *(Continued)*

Parameter	Description
	<i>device</i>) at the end of its name, and the status refers to the Online status. Based on your requirement, you can select or deselect Online or Offline. - Interface types of the Test Agent: - Normal—Network interfaces that are used for testing traffic. - Management—Network interfaces used by the Test Agents to connect to Routing Director. - Unknown—Network interfaces that the Test Agent is unable to identify or that have been removed. Based on your requirement, you can select or deselect Normal, Management, or Unknown. You can enable the Display Device Names toggle button to display the hostnames of the routers that are associated with the Test Agents. You can also enable the Hide unsupported toggle button to hide devices that are not supported by the plugin you have selected. To search for a specific interface or a Test Agent, enter one or more keywords in the search text box. You can search based on the name, description, IP and MAC address of the interfaces or the Test Agents. You can also search based on the device name, device model, and device MAC addresses if searching for a Test Agent associated with the device. You can hover over an interface name to view the details such as the Test Agent name, bind interface, bind family, and to check if the Test Agent is run on a device or not.

Table 265: DNS Parameters *(Continued)*

Parameter	Description
	You can also remove the interface filters you have selected and reset the filters to the default mode by using the Reset filters button.
Lookup names*	Specify the domain name to lookup. For example, example.net Lookup is the process of querying the server to translate a domain name into IP address. When you click the Lookup names text box, a Select Lookup name page appears. You can add a name in the Lookup text box. To add more than one Lookup names, click + Add Lookup.
Time between requests	Specify the time that Test Agent should wait between successive DNS queries. Unit—seconds (s) Default value—10.00 s Range—0.01 s through 3600 s
DNS server	Specify the IP address of the DNS server to query. If left empty, Test Agents use the default DNS server. Maximum Length—200 characters

Table 265: DNS Parameters *(Continued)*

Parameter	Description
DNS record type	Select the DNS record type. A DNS record is a set of unstructured data stored in a DNS database. The database consists of information on a domain and its services. DNS has different resource records. Each record type has different functions in the resolution process. • A (Address)—Associates a domain name with an IPv4 address. • AAAA (IPv6 Address)—Associates a domain name with an IPv6 address. • SOA Record (Start of Authority)—Provides information about a DNS zone. • MX (Mail Exchange)—Associates a domain name to the email server responsible for receiving an email. • NS (Name Server)—Provides domain information of an address. It specifies the servers responsible for hosting DNS servers for a particular domain. • TXT (Text)—Provides storage for text data. It is also used for domain verification. • PTR (Pointer)—Performs reverse DNS operation. It associates an IP address with a domain name. • CNAME (Canonical Name)—Maps domain names to another domain and not an IP address. By default, the record type is A.
Thresholds for errored seconds (ES)	

Table 265: DNS Parameters *(Continued)*

Parameter	Description
Timeout	Specify the timeout value. If no response is received for the DNS request within the defined timeout period, an errored-second is raised unless the request lifetime expires. Unit—Milliseconds (ms) Default value—1000 ms Range—1ms through 30000 ms
Advanced	
Request lifetime	Specify the timeout value. If no response is received for the DNS request within the defined timeout period, an errored-second is raised unless the request lifetime expires. Request lifetime value is the duration for which a DNS request is alive. It determines how long a request persists without getting terminated. Unit—Milliseconds (ms) Default value—5000 ms Range—1 ms through 30000 ms
Recursive requests	Enable or disable the Recursive Requests toggle button. Recursive Request is the process where a DNS server queries other DNS servers until it resolves the lookup address. If you enable the Recursive Requests toggle button, the Recursion Desired flag is set on the DNS queries. By default, the toggle button is enabled.

Table 265: DNS Parameters *(Continued)*

Parameter	Description
Response code	Select the DNS response code. You can select one of the following response codes: • NOERROR—Indicates that the query was successful. • SERVFAIL—Indicates that the server has failed to complete the request. • NXDOMAIN—Indicates that the domain name does not exist. • REFUSED—Indicates that the server refused to perform the operation. • NOTAUTH—Indicates that the server is not authoritative for the zone. The DNS response code can be any value from 0 to 9. The value indicates the outcome of a DNS query. For more information on DNS response codes and strings with descriptions, see IANA link.
Expected response	Specify an expected response from the DNS server. If the actual response does not match the expected response, an errored-second is raised. In case, the response consists of more than one codes, at least one code must match the expected response. For more information on DNS response codes and strings with descriptions, see IANA link.

Table 266: HTTP Parameters

Parameter	Description
General	

Table 266: HTTP Parameters *(Continued)*

Parameter	Description
Clients*	Select one or more Test Agent interfaces on which you want to run the Test. To select one or more Test Agent interfaces: 1. Click the Clients text box. The Test Agent/Interfaces page appears. 2. Select one or more Test Agent interfaces. For information on filtering Test Agent interfaces, see "Search for Test Agent interfaces by using Filters" on page 972. 3. Click OK. NOTE: On the Test Agent/Interfaces page, you can filter the Test Agent interfaces based on the following criteria: • Interface name or the Test Agent name, or both. Based on your requirement, you can select or deselect Test Agents or Interfaces. • IPv4 address (with or without IPv4 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv4 or No IPv4 to view or hide the interfaces that are associated with the IPv4 family and the interfaces that are not associated with the IPv4 family. • IPv6 address (with or without IPv6 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv6 or No IPv6 to view or hide the interfaces that are associated with the IPv6 family and the interfaces that are not associated with the IPv6 family.

Table 266: HTTP Parameters *(Continued)*

Parameter	Description
	• Deployment type of a Test Agent: • Application—Filter out the Test Agents deployed in a device as an application. • Device—Filter out the Test Agents associated with a device. Based on your requirement, you can select or deselect Application or Device to view or hide the Test Agents deployed in a device as an application or the Test Agents associated with a device. • Connection types between the interface and the router: • Link—Indicates that the status is Up and there is an active connection for the interface. • No Link—Indicates that the status is Down and there is no active connection for the interface. Based on your requirement, you can select or deselect Link or No Link. • Connection statuses of the Test Agent: • Online—The Test Agent is connected to Routing Director. • Offline—The Test Agent is not connected to Routing Director. When the interface supports RPM plugins, the interface contains <i>(On device)</i> at the end of its name indicating the Device Online status.

Table 266: HTTP Parameters *(Continued)*

Parameter	Description
	When the interface supports only the normal plugins, the interface will not display <i>(On device)</i> at the end of its name, and the status refers to the Online status. Based on your requirement, you can select or deselect Online or Offline. - Interface types of the Test Agent: - Normal—Network interfaces that are used for testing traffic. - Management—Network interfaces used by the Test Agents to connect to Routing Director. - Unknown—Network interfaces that the Test Agent is unable to identify or that have been removed. Based on your requirement, you can select or deselect Normal, Management, or Unknown. You can enable the Display Device Names toggle button to display the hostnames of the routers that are associated with the Test Agents. You can also enable the Hide unsupported toggle button to hide devices that are not supported by the plugin you have selected. To search for a specific interface or a Test Agent, enter one or more keywords in the search text box. You can search based on the name, description, IP and MAC address of the interfaces or the Test Agents. You can also search based on the device name, device model, and device MAC addresses if searching for a Test Agent associated with the device. You can also hover over an interface name to view the details such as the Test Agent name, bind

Table 266: HTTP Parameters *(Continued)*

Parameter	Description
	interface, bind family, and to check if the Test Agent is run on a device or not. You can also remove the interface filters you have selected and reset the filters to the default mode by using the Reset filters button.
URLs*	Specify the URL for the HTTP request. Maximum Length—200 characters
Time between requests	Specify the time taken between successive HTTP requests initiated by Test Agent. Unit—Seconds (s) Default value—10.00 s Range—0.01 s through 3600 s
Thresholds for errored seconds (ES)	
Timeout	Specify the timeout value. If no response is received to the HTTP request within the defined timeout period, an errored second for timeout is raised. Unit—Milliseconds (ms) Default value—1000 ms Range—1 ms through 30000 ms
Response content	Specify the regular expression (case-insensitive) against which the HTTP response content is validated. Maximum Length—50 characters
Advanced	

Table 266: HTTP Parameters *(Continued)*

Parameter	Description
Request lifetime	Specify the maximum time to wait for a response before the HTTP request is canceled. If this time expires without a response, an errored second for lifetime expiry is raised. Unit—Milliseconds (ms) Default value—5000 ms Range—1 ms through 30000 ms
HTTP response code	Specify the HTTP response code for Test Agent. If the HTTP response code does not match the specified response code, an errored-second is raised. For more information on HTTP response codes, see RFC 9110.
Proxy server	Specify the IP address of the HTTP proxy server optionally.
Proxy port	Specify the HTTP proxy port when using HTTP proxy server. Default value—8080 Range—1 through 65535

Table 266: HTTP Parameters *(Continued)*

Parameter	Description
Proxy authentication	Select the authentication method that the Test Agent uses when connecting to a proxy server. Select one of the following authentication method: • None—Indicates that you can access the remote endpoint without any authentication. • Basic—Indicates that you can access the remote endpoint by using a username and password. • Digest—Indicates that you can access the remote endpoint by using a username and password, but it will send a hashed version of the password, making the password resistant to attacks. • Ntlm—Indicates that you can access a remote endpoint by using Single Sign-on. That is, without a password. Proxy authenticates the incoming request from a Test Agent. This ensures that only the authorized users have access to the internet.
Proxy username	Specify the username for authorized access to a proxy.
Proxy password	Specify the password for authorized access to a proxy server.

Table 267: Ping Parameters

Parameter	Description
General	

Table 267: Ping Parameters *(Continued)*

Parameter	Description
Clients*	Select one or more Test Agent interfaces on which you want to run the Test. To select one or more Test Agent interfaces: 1. Click the Clients text box. The Select interfaces page appears. 2. Select one or more Test Agent interfaces. For information on filtering Test Agent interfaces, see "Search for Test Agent interfaces by using Filters" on page 978. 3. Click OK. NOTE: On the Select interfaces page, you can filter the Test Agent interfaces based on the following criteria:

Table 267: Ping Parameters (Continued)

Parameter	Description
	- Interface name or the Test Agent name, or both. Based on your requirement, you can select or deselect Test Agents or Interfaces. - IPv4 address (with or without IPv4 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv4 or No IPv4 to view or hide the interfaces that are associated with the IPv4 family and the interfaces that are not associated with the IPv4 family. - IPv6 address (with or without IPv6 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv6 or No IPv6 to view or hide the interfaces that are associated with the IPv6 family and the interfaces that are not associated with the IPv6 family. - Deployment type of a Test Agent: - Application—Filter out the Test Agents deployed in a device as an application. - Device—Filter out the Test Agents associated with a device. Based on your requirement, you can select or deselect Application or Device to view or hide the Test Agents deployed in a device as an application or the Test Agents associated with a device. - Connection types between the interface and the router:

Table 267: Ping Parameters *(Continued)*

Parameter	Description
	• Link—Indicates that the status is Up and there is an active connection for the interface. • No Link—Indicates that the status is Down and there is no active connection for the interface. Based on your requirement, you can select or deselect Link or No Link. • Connection statuses of the Test Agent: • Online—The Test Agent is connected to Routing Director. • Offline—The Test Agent is not connected to Routing Director. When the interface supports RPM plugins, the interface contains <i>(On device)</i> at the end of its name indicating the Device Online status. When the interface supports only the normal plugins, the interface will not display <i>(On device)</i> at the end of its name, and the status refers to the Online status. Based on your requirement, you can select or deselect Online or Offline. • Interface types of the Test Agent: • Normal—Network interfaces that are used for testing traffic. • Management—Network interfaces used by the Test Agents to connect to Routing Director.

Table 267: Ping Parameters *(Continued)*

Parameter	Description
	Unknown—Network interfaces that the Test Agent is unable to identify or that have been removed. Based on your requirement, you can select or deselect Normal, Management, or Unknown. You can enable the Display Device Names toggle button to display the hostnames of the routers that are associated with the Test Agents. You can also enable the Hide unsupported toggle button to hide devices that are not supported by the plugin you have selected. To search for a specific interface or a Test Agent, enter one or more keywords in the search text box. You can search based on the name, description, IP and MAC address of the interfaces or the Test Agents. You can also search based on the device name, device model, and device MAC addresses if searching for a Test Agent associated with the device. You can also hover over an interface name to view the details such as the Test Agent name, bind interface, bind family, and to check if the Test Agent is run on a device or not. You can also remove the interface filters you have selected and reset the filters to the default mode by using the Reset filters button.

Table 267: Ping Parameters (Continued)

Parameter	Description
Hosts*	Specify the hostname or the destination IP. A host is the remote endpoint to which the Test Agent sends the request. When you click the Host text box, the Select Host page appears where you can enter hostnames. To add more than one hosts, click + Add Host and specify the following: • Host—The hostname or the IP address of the remote endpoint. Maximum length—255 characters • Name—The text box is automatically populated based on the data you specified in the Host text box.
Time between requests	Specify the time taken between successive ping requests initiated by Test Agent. Unit—Seconds (s) Default value—10.00 s Range—0.01 s to 3600 s
Thresholds for error seconds (ES)	
Delay	Specify the maximum threshold for the delay. If the round-trip time for a Ping response exceeds this value, an errored-second is raised. Unit—Milliseconds (ms) Default value—1000 ms Range—1 ms through 30000 ms

Table 267: Ping Parameters *(Continued)*

Parameter	Description
DV (Delay variance)	Specify the maximum threshold value for jitter or delay variance (DV). DV is the difference between the maximum and minimum delay experienced in a one second metrics interval. If the DV exceeds the defined value, an errored-second is raised. Delay variation (DV) occurs when different packets take different amount of time to travel from a Test Agent to a remote endpoint. Packets are sent at regular interval of time and if variation is experienced in consecutive packets, the Test Agent generates an errored-second event. Unit—Milliseconds (ms) Default value—500 ms Range—0 ms through 10000 ms
Advanced	
UDP echo	Enable UDP Echo protocol instead of ICMP Ping. By default, the toggle button is disabled.
Payload	Specify the size of the ping payload. Payload is the actual data in a request packet. Unit—Bytes Default value—56 bytes Range—0 byte through 65000 bytes

Table 267: Ping Parameters *(Continued)*

Parameter	Description
TTL (Time to Live)	Specify the number of times the packets hop before a device discards the packet. Time to live (TTL) value indicates lifespan of a request packet. TTL prevents infinite loop in a network when the packet fails to reach the destination. For example, if you have entered the TTL value as 64, every time the packet passes a device, the value is reduced by one until the packet reaches the remote server. If the entered value reaches zero before reaching the remote endpoint, the packet is discarded. Default value—64 Range—1 through 255
Request Lifetime	Specify the request lifetime value. Request lifetime value is the maximum duration for which a ping request is alive. It determines how long a request persists before it is counted as lost. Unit—Milliseconds (ms) Default value—2000 ms Range—1 ms through 30000 ms

Table 267: Ping Parameters *(Continued)*

Parameter	Description
DSCP/IPP	Specify the Differentiated Services Code Point (DSCP) or the IP Precedence (IPP) value that is used in the IP packet headers. The IPP is the three-bit binary values (Precedence) in the ToS field of the IP header. An IPP value can be in the 0-7 range. IPP value informs the router about the priority of the packet. The higher the IPP value, the more the priority of the packet. See RFC 791 for more information. DSCP is a six-bit binary value in the DS field of the IP header. DSCP value facilitates QoS for traffic management through the Best Effort, Assured Forwarding, Class Selector, and the Expedited Forwarding categories. See RFC 2474 for more information. Default value—0 Range—0 through 255

Table 268: RPM CISCO TWAMP Reflector

Parameter	Description
General	

Table 268: RPM CISCO TWAMP Reflector (Continued)

Parameter	Description
Reflectors*	Select one or more Test Agent interfaces on which you want to run the Test. To select one or more Test Agent interfaces: 1. Click the Clients text box. The Test Agent/Interfaces choices page appears. 2. Select one or more Test Agent interfaces. For information on filtering Test Agent interfaces, see "Search for Test Agent interfaces by using Filters" on page 986. 3. Click OK. NOTE: On the Test Agent/Interfaces choices page, you can filter the Test Agent interfaces based on the following criteria: • Interface name or the Test Agent name, or both. Based on your requirement, you can select or deselect Test Agents or Interfaces. • IPv4 address (with or without IPv4 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv4 or No IPv4 to view or hide the interfaces that are associated with the IPv4 family and the interfaces that are not associated with the IPv4 family. • IPv6 address (with or without IPv6 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv6 or No IPv6 to view or hide the interfaces that are associated with the IPv6 family and the interfaces that are not associated with the IPv6 family.

Table 268: RPM CISCO TWAMP Reflector *(Continued)*

Parameter	Description
	- Deployment type of a Test Agent: - Application—Filter out the Test Agents deployed in a device as an application. - Device—Filter out the Test Agents associated with a device. Based on your requirement, you can select or deselect Application or Device to view or hide the Test Agents deployed in a device as an application or the Test Agents associated with a device. - Connection types between the interface and the router: - Link—Indicates that the status is Up and there is an active connection for the interface. - No Link—Indicates that the status is Down and there is no active connection for the interface. Based on your requirement, you can select or deselect Link or No Link. - Connection statuses of the Test Agent: - Online—The Test Agent is connected to Routing Director. - Offline—The Test Agent is not connected to Routing Director. When the interface supports RPM plug-ins, the interface contains <i>(On device)</i> at the end of its name indicating the Device Online status.

Table 268: RPM CISCO TWAMP Reflector (Continued)

Parameter	Description
	When the interface supports only the normal plug ins, the interface will not display <i>(On device)</i> at the end of its name, and the status refers to the Online status. Based on your requirement, you can select or deselect Online or Offline. - Interface types of the Test Agent: - Normal—Network interfaces that are used for testing traffic. - Management—Network interfaces used by the Test Agents to connect to Routing Director. - Unknown—Network interfaces that the Test Agent is unable to identify or that have been removed. Based on your requirement, you can select or deselect Normal, Management, or Unknown. You can enable the Display Device Names toggle button to display the hostnames of the routers that are associated with the Test Agents. You can also enable the Hide unsupported toggle button to hide devices that are not supported by the plug-in you have selected. To search for a specific interface or a Test Agent, enter one or more keywords in the search text box. You can search based on the name, description, IP and MAC address of the interfaces or the Test Agents. You can also search based on the device name, device model, and device MAC addresses if searching for a Test Agent associated with the device. You can also hover over an interface name to view the details such as the Test Agent name, bind

Table 268: RPM CISCO TWAMP Reflector *(Continued)*

Parameter	Description
	interface, bind family, and to check if the Test Agent is run on a device or not. You can also remove the interface filters you have selected and reset the filters to the default mode by using the Reset filters button.
Ports*	Specify the Test session port value. When you click the text box, the Select Test Session Port Value page appears where you can add port values. On this page: • Test Session Port Value—Specify the local port used for the test session. Maximum length—65535 characters Minimum length—1 Default length—862 • Host—Specify IPv4 or IPv6 Address of TWAMP client. If no address is specified, Routing Director accepts connections from any TWAMP sender or client. Maximum length—255 characters • Client Test session port—Specify the port used by the client for the test session. If no port value is specified, Routing Director accepts connections from any port of the TWAMP client. Maximum length—65535 characters Minimum length—1 • Name—This text box is automatically populated based on the data you specified in the Host text box.
Advanced	

Table 268: RPM CISCO TWAMP Reflector *(Continued)*

Parameter	Description
Collection interval*	Specify the collection interval. Collection interval is the frequency at which the results are collected from the device. Collection interval must be larger than the value specified in the Timeout period field. Unit—Seconds (s) Default value—30s Range—30s through 300s
Session Timeout	Specify the duration of time after which an idle session is automatically ended. Unit—Seconds (s) Default value—35s Range— 10s through 300s
Timeout period	Specify the inactivity timeout period for the test session. The session terminates if no activity happens within the specified time. Unit—Seconds (s) Range—1s through 604800s
Inactivity-timer value	Specify the inactivity timeout period for the TWAMP server. The session terminates if no activity happens within the specified time. Unit—Seconds (s) Range—1s through 6000s
Server Test Session Port IP	Specify the IPv4 or IPv6 address of the TWAMP reflector used for test sessions. Maximum length—255 characters

Table 268: RPM CISCO TWAMP Reflector *(Continued)*

Parameter	Description
Control Session Port	Specify the port used for the control session. Specify 0 for TWAMP Light and specify any value starting from 862 through 65535 for managed TWAMP. Default value—0 Range—0 through 65535
Session Duration	Specify the total duration the test session runs. Unit—Seconds (s) Default value—Unlimited (the session continues until the session is manually stopped.) Range— 60s through 86400s
VRF	Specify the VRF (Virtual Routing and Forwarding) instance to use for the test session. Default value—Any Maximum Length—255 characters
Session ID	Specify the session ID used for the TWAMP light test session. Default value—1 Range—1 through 65535

Table 269: IPTV MPEG Parameters

Parameter	Description
General	

Table 269: IPTV MPEG Parameters *(Continued)*

Parameter	Description
Clients*	Select one or more Test Agent interfaces on which you want to run the Test. To select one or more Test Agent interfaces: 1. Click the Clients text box. The Test Agent/Interfaces choices page appears. 2. Select one or more Test Agent interfaces. For information on filtering Test Agent interfaces, see "Search for Test Agent interfaces by using Filters." on page 992 3. Click OK. NOTE: On the Test Agent/Interfaces choices page, you can filter the Test Agent interfaces based on the following criteria: • Interface name or the Test Agent name, or both. Based on your requirement, you can select or deselect Test Agents or Interfaces. • IPv4 address (with or without IPv4 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv4 or No IPv4 to view or hide the interfaces that are associated with the IPv4 family and the interfaces that are not associated with the IPv4 family. • IPv6 address (with or without IPv6 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv6 or No IPv6 to view or hide the interfaces that are associated with the IPv6 family and the interfaces that are not associated with the IPv6 family.

Table 269: IPTV MPEG Parameters *(Continued)*

Parameter	Description
	• Deployment type of a Test Agent: • Application—Filter out the Test Agents deployed in a device as an application. • Device—Filter out the Test Agents associated with a device. Based on your requirement, you can select or deselect Application or Device to view or hide the Test Agents deployed in a device as an application or the Test Agents associated with a device. • Connection types between the interface and the router: • Link—Indicates that the status is Up and there is an active connection for the interface. • No Link—Indicates that the status is Down and there is no active connection for the interface. Based on your requirement, you can select or deselect Link or No Link. • Connection statuses of the Test Agent: • Online—The Test Agent is connected to Routing Director. • Offline—The Test Agent is not connected to Routing Director. When the interface supports RPM plugins, the interface contains <i>(On device)</i> at the end of its name indicating the Device Online status.

Table 269: IPTV MPEG Parameters (Continued)

Parameter	Description
	When the interface supports only the normal plugins, the interface will not display (<i>On device</i>) at the end of its name, and the status refers to the Online status. Based on your requirement, you can select or deselect Online or Offline. - Interface types of the Test Agent: - Normal—Network interfaces that are used for testing traffic. - Management—Network interfaces used by the Test Agents to connect to Routing Director. - Unknown—Network interfaces that the Test Agent is unable to identify or that have been removed. Based on your requirement, you can select or deselect Normal, Management, or Unknown. You can enable the Display Device Names toggle button to display the hostnames of the routers that are associated with the Test Agents. You can also enable the Hide unsupported toggle button to hide devices that are not supported by the plugin you have selected. To search for a specific interface or a Test Agent, enter one or more keywords in the search text box. You can search based on the name, description, IP and MAC address of the interfaces or the Test Agents. You can also search based on the device name, device model, and device MAC addresses if searching for a Test Agent associated with the device. You can also hover over an interface name to view the details such as the Test Agent name, bind

Table 269: IPTV MPEG Parameters (Continued)

Parameter	Description
	interface, bind family, and to check if the Test Agent is run on a device or not. You can also remove the interface filters you have selected and reset the filters to the default mode by using the Reset filters button.
IPTV Channels*	Specify one or more IPTV channels to monitor from the preset IPTV channel list. When you click IPTV Channels, the Select Channel page appears where you can enter channels. To add more than one channels, click + Add channel and specify the following: • Multicast group address—Specify the IPv4 or IPv6 multicast address of the IPTV channel. Maximum length—64 characters. • Multicast group port—Specify the UDP destination port used in the IPTV multicast stream. Range—1 through 65535. • Multicast group source—Specify the source address of the multicast group. When provided, IGMPv3 (for IPv4) or MLDv2 (for IPv6) is used for joining the channels. • Program number—Specify the program number of the channel within a multicast group. • Name—Specify the name of the IPTV channel.
Thresholds for error seconds (ES)	

Table 269: IPTV MPEG Parameters *(Continued)*

Parameter	Description
MPEG loss (CC errors/s)	Specify the maximum allowed MPEG packet loss (Continuity Count errors) per second. An errored-second is raised if the MPEG loss exceeds this threshold. Continuity Counter (CC) error indicates an issue with the packet sequence in an MPEG transport stream.
PAT/PMT delay variation (ms)	Specify the maximum tolerated jitter (delay variation) of the received MPEG stream. An errored-second is triggered if the PCR jitter or RTP jitter exceeds this threshold. PCR provides timestamps to maintain synchronization of audio and video, RTP provides end-to-end delivery of real-time data over IP networks.
PAT/PMT interval (s)	Specify the maximum allowable time interval between Program Allocation Table (PAT)/ Program Map Table (PMT) transmissions. PAT has information on all the available programs in a transport stream. PAT maps all the programs numbers to PMT packet identifiers. PMT has information on the streams and lists the PIDs of all the streams associated with a program.
PID interval (s)	Specify the maximum allowable time interval between audio or video Packet Identifiers (PIDs) as specified by Program Map Table (PMT). PID is a 13-bit field in the header of a transport stream used to identify different data streams, such as video, audio, subtitles, within a transport stream.

Table 270: Netflix Speedtest Parameters

Parameter	Description
General	

Table 270: Netflix Speedtest Parameters *(Continued)*

Parameter	Description
Clients*	Select one or more Test Agent interfaces on which you want to run the Test. To select one or more Test Agent interfaces: 1. Click the Clients text box. The Test Agent/Interfaces choices page appears. 2. Select one or more Test Agent interfaces. For information on filtering Test Agent interfaces, see "Search for Test Agent interfaces by using Filters" on page 998. 3. Click OK. NOTE: On the Test Agent/Interfaces choices page, you can filter the Test Agent interfaces based on the following criteria: • Interface name or the Test Agent name, or both. Based on your requirement, you can select or deselect Test Agents or Interfaces. • IPv4 address (with or without IPv4 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv4 or No IPv4 to view or hide the interfaces that are associated with the IPv4 family and the interfaces that are not associated with the IPv4 family. • IPv6 address (with or without IPv6 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv6 or No IPv6 to view or hide the interfaces that are associated with the IPv6 family and the interfaces that are not associated with the IPv6 family.

Table 270: Netflix Speedtest Parameters *(Continued)*

Parameter	Description
	- Deployment type of a Test Agent: - Application—Filter out the Test Agents deployed in a device as an application. - Device—Filter out the Test Agents associated with a device. Based on your requirement, you can select or deselect Application or Device to view or hide the Test Agents deployed in a device as an application or the Test Agents associated with a device. - Connection types between the interface and the router: - Link—Indicates that the status is Up and there is an active connection for the interface. - No Link—Indicates that the status is Down and there is no active connection for the interface. Based on your requirement, you can select or deselect Link or No Link. - Connection statuses of the Test Agent: - Online—The Test Agent is connected to Routing Director. - Offline—The Test Agent is not connected to Routing Director. When the interface supports RPM plugins, the interface contains <i>(On device)</i> at the end of its name indicating the Device Online status.

Table 270: Netflix Speedtest Parameters *(Continued)*

Parameter	Description
	When the interface supports only the normal plugins, the interface will not display <i>(On device)</i> at the end of its name, and the status refers to the Online status. Based on your requirement, you can select or deselect Online or Offline. - Interface types of the Test Agent: - Normal—Network interfaces that are used for testing traffic. - Management—Network interfaces used by the Test Agents to connect to Routing Director. - Unknown—Network interfaces that the Test Agent is unable to identify or that have been removed. Based on your requirement, you can select or deselect Normal, Management, or Unknown. You can enable the Display Device Names toggle button to display the hostnames of the routers that are associated with the Test Agents. You can also enable the Hide unsupported toggle button to hide devices that are not supported by the plugin you have selected. To search for a specific interface or a Test Agent, enter one or more keywords in the search text box. You can search based on the name, description, IP and MAC address of the interfaces or the Test Agents. You can also search based on the device name, device model, and device MAC addresses if searching for a Test Agent associated with the device. You can also hover over an interface name to view the details such as the Test Agent name, bind

Table 270: Netflix Speedtest Parameters *(Continued)*

Parameter	Description
	interface, bind family, and to check if the Test Agent is run on a device or not. You can also remove the interface filters you have selected and reset the filters to the default mode by using the Reset filters button.
Measurement Cycle period	Specify the time interval between successive measurement cycles. On each measurement cycle, the Test Agent downloads Netflix test segments over HTTPS from one or several OCAs (Open Connect Appliances) operated by Netflix. OCA is a server appliance to cache and deliver content locally to subscribers, reducing the load on Netflix's central servers. Unit—Minutes (min) Default value—15 min Range—1 min through 60 min
Upload or Download	Select the type of bandwidth to measure the rate of the data transfer. Upload data transfer is calculated from the Test Agent to OCAs, while download is calculated from the OCAs to the Test Agent.
Advanced	
Minimum Duration (s)	Specify the minimum duration of speedtest that must be run before checking for throughput stabilization. Unit—Seconds (s) Default value—2 s Range—1 s through 3600 s

Table 270: Netflix Speedtest Parameters *(Continued)*

Parameter	Description
Maximum Duration (s)	Specify the maximum duration of speedtest that must be run to check for throughput stabilization. Unit—Seconds (s) Default value—30 s Range—5 s through 3600 s
Minimum Concurrent OCA connections	Specify the minimum number of concurrent Open Connect Appliance (OCA) connections used for bandwidth testing. Default value—3. Range—1 through 5.
Maximum Concurrent OCA connections	Specify the maximum number of concurrent Open Connect Appliance (OCA) connections used for bandwidth testing. Default value—5 Range—1 through 5
Thresholds for error seconds (ES)	
Minimum bandwidth (Mbit/s)	Specify the minimum expected bandwidth during the speedtest. An errored-second is raised if the bandwidth drops below the defined value.
Maximum latency (ms)	Specify the minimum expected bandwidth during the speedtest. An errored-second is raised if the bandwidth drops below the defined value.

Table 271: OTT - HLS Parameters

Parameter	Description
General	

Table 271: OTT - HLS Parameters (Continued)

Parameter	Description
Clients*	Select one or more Test Agent interfaces on which you want to run the Test. To select one or more Test Agent interfaces: 1. Click the Clients text box. The Test Agent/Interfaces choices page appears. 2. Select one or more Test Agent interfaces. For information on filtering Test Agent interfaces, see "Search for Test Agent interfaces by using Filters" on page 1004. 3. Click OK. NOTE: On the Test Agent/Interfaces choices page, you can filter the Test Agent interfaces based on the following criteria: • Interface name or the Test Agent name, or both. Based on your requirement, you can select or deselect Test Agents or Interfaces. • IPv4 address (with or without IPv4 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv4 or No IPv4 to view or hide the interfaces that are associated with the IPv4 family and the interfaces that are not associated with the IPv4 family. • IPv6 address (with or without IPv6 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv6 or No IPv6 to view or hide the interfaces that are associated with the IPv6 family and the interfaces that are not associated with the IPv6 family.

Table 271: OTT - HLS Parameters *(Continued)*

Parameter	Description
	• Deployment type of a Test Agent: • Application—Filter out the Test Agents deployed in a device as an application. • Device—Filter out the Test Agents associated with a device. Based on your requirement, you can select or deselect Application or Device to view or hide the Test Agents deployed in a device as an application or the Test Agents associated with a device. • Connection types between the interface and the router: • Link—Indicates that the status is Up and there is an active connection for the interface. • No Link—Indicates that the status is Down and there is no active connection for the interface. Based on your requirement, you can select or deselect Link or No Link. • Connection statuses of the Test Agent: • Online—The Test Agent is connected to Routing Director. • Offline—The Test Agent is not connected to Routing Director. When the interface supports RPM plugins, the interface contains <i>(On device)</i> at the end of its name indicating the Device Online status.

Table 271: OTT - HLS Parameters (Continued)

Parameter	Description
	When the interface supports only the normal plugins, the interface will not display (<i>On device</i>) at the end of its name, and the status refers to the Online status. Based on your requirement, you can select or deselect Online or Offline. - Interface types of the Test Agent: - Normal—Network interfaces that are used for testing traffic. - Management—Network interfaces used by the Test Agents to connect to Routing Director. - Unknown—Network interfaces that the Test Agent is unable to identify or that have been removed. Based on your requirement, you can select or deselect Normal, Management, or Unknown. You can enable the Display Device Names toggle button to display the hostnames of the routers that are associated with the Test Agents. You can also enable the Hide unsupported toggle button to hide devices that are not supported by the plugin you have selected. To search for a specific interface or a Test Agent, enter one or more keywords in the search text box. You can search based on the name, description, IP and MAC address of the interfaces or the Test Agents. You can also search based on the device name, device model, and device MAC addresses if searching for a Test Agent associated with the device. You can also hover over an interface name to view the details such as the Test Agent name, bind

Table 271: OTT - HLS Parameters (Continued)

Parameter	Description
	interface, bind family, and to check if the Test Agent is run on a device or not. You can also remove the interface filters you have selected and reset the filters to the default mode by using the Reset filters button.
URLs*	Specify the URL of the video stream to which you send the connection requests. The URL can either be a playlist linking to other playlists for the different variants, or a playlist containing the segments. The file extension is .m3u or .m3u8. When you click URLs, the Select URL page appears where you can enter URLs. To add more than one URLs, click + Add URL and specify the following: • URL—The host or the IP address of the remote endpoint. • Name—The text box is automatically populated based on the data you specified in URL.
Buffer size (s)	Specify the target duration of the buffered data. If the duration of the buffer data drops below the defined value, new segments are downloaded. Unit—Seconds (ms) Default value—60 s Minimum—1 s

Table 271: OTT - HLS Parameters (*Continued*)

Parameter	Description
Initial Buffering (s)	Specify the duration of the initial buffered data required before the playback begins. Unit—Seconds (ms) Default value—10 s Minimum—1 s
Loop	Enable or disable the Loop button. If enabled, the playback loops when the stream reaches the end of the playlist. If disabled, the playback stops, and errored-seconds are raised. By default, the toggle button is enabled.
Thresholds for error seconds (ES)	
Minimum playback rate (Mbit/s)	Specify the minimum expected playback rate of the video stream. An errored second is triggered if its value drops below this threshold. Unit—Megabit per second (Mbit/s) Minimum value—0 Mbit/s
Minimum download rate (Mbit/s)	Specify the minimum expected playback rate of the video stream. An errored-second is raised if the download rate drops below the defined value. Unit—Megabit per second (Mbit/s) Minimum value—0 Mbit/s

Table 271: OTT - HLS Parameters *(Continued)*

Parameter	Description
Minimum selected rate (Mbit/s)	Specify the minimum expected download rate for the segments. An errored-second is raised if the selected rate drops below the defined value. Variants are the versions of media streams that is encoded at a specific rate. They enable adaptive streaming where the Test Agents can select variants to provide the best experience. Unit—Megabit per second (Mbit/s) Minimum value—0 Mbit/s
Minimum buffer size (s)	Specify the minimum expected duration of the data segment buffered in the Test Agent. An errored-second is triggered if its value drops below the defined value. Unit—Megabit per second (Mbit/s) Minimum value—0 Mbit/s

Table 272: TCP Parameters

Parameter	Description
General	

Table 272: TCP Parameters (Continued)

Parameter	Description
Clients*	Select one or more Test Agent interfaces which you want to use as TCP clients. To select one or more Test Agent interfaces: 1. Click Clients. The Test Agent/Interfaces choices page appears. 2. Select one or more Test Agent interfaces. For information on filtering Test Agent interfaces, see "Search for Test Agent interfaces by using Filters" on page 1010. 3. Click OK. NOTE: On the Test Agent/Interfaces choices page, you can filter the Test Agent interfaces based on the following criteria: • Interface name or the Test Agent name, or both. Based on your requirement, you can select or deselect Test Agents or Interfaces. • IPv4 address (with or without IPv4 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv4 or No IPv4 to view or hide the interfaces that are associated with the IPv4 family and the interfaces that are not associated with the IPv4 family. • IPv6 address (with or without IPv6 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv6 or No IPv6 to view or hide the interfaces that are associated with the IPv6 family and the interfaces that are not associated with the IPv6 family.

Table 272: TCP Parameters (Continued)

Parameter	Description
	• Deployment type of a Test Agent: • Application—Filter out the Test Agents deployed in a device as an application. • Device—Filter out the Test Agents associated with a device. Based on your requirement, you can select or deselect Application or Device to view or hide the Test Agents deployed in a device as an application or the Test Agents associated with a device. • Connection types between the interface and the router: • Link—Indicates that the status is Up and there is an active connection for the interface. • No Link—Indicates that the status is Down and there is no active connection for the interface. Based on your requirement, you can select or deselect Link or No Link. • Connection statuses of the Test Agent: • Online—The Test Agent is connected to Routing Director. • Offline—The Test Agent is not connected to Routing Director. When the interface supports RPM plugins, the interface contains <i>(On device)</i> at the end of its name indicating the Device Online status.

Table 272: TCP Parameters *(Continued)*

Parameter	Description
	When the interface supports only the normal plugins, the interface will not display <i>(On device)</i> at the end of its name, and the status refers to the Online status. Based on your requirement, you can select or deselect Online or Offline. - Interface types of the Test Agent: - Normal—Network interfaces that are used for testing traffic. - Management—Network interfaces used by the Test Agents to connect to Routing Director. - Unknown—Network interfaces that the Test Agent is unable to identify or that have been removed. Based on your requirement, you can select or deselect Normal, Management, or Unknown. You can enable the Display Device Names toggle button to display the hostnames of the routers that are associated with the Test Agents. You can also enable the Hide unsupported toggle button to hide devices that are not supported by the plugin you have selected. To search for a specific interface or a Test Agent, enter one or more keywords in the search text box. You can search based on the name, description, IP and MAC address of the interfaces or the Test Agents. You can also search based on the device name, device model, and device MAC addresses if searching for a Test Agent associated with the device. You can also hover over an interface name to view the details such as the Test Agent name, bind

Table 272: TCP Parameters *(Continued)*

Parameter	Description
	interface, bind family, and to check if the Test Agent is run on a device or not. You can also remove the interface filters you have selected and reset the filters to the default mode by using the Reset filters button.

Table 272: TCP Parameters (Continued)

Parameter	Description
Server*	Select one or more Test Agent interfaces that you want to use as UDP servers. To select one or more Test Agent interfaces: 1. Click Clients. The Test Agent/Interfaces choices page appears. 2. Select one or more Test Agent interfaces. For information on filtering Test Agents, see "Search for Test Agent interfaces by using Filters." on page 1014 3. Click OK. NOTE: On the Test Agent/Interfaces choices page, you can filter the Test Agent interfaces based on the following criteria: • Interface name or the Test Agent name, or both. Based on your requirement, you can select or deselect Test Agents or Interfaces. • IPv4 address (with or without IPv4 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv4 or No IPv4 to view or hide the interfaces that are associated with the IPv4 family and the interfaces that are not associated with the IPv4 family. • IPv6 address (with or without IPv6 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv6 or No IPv6 to view or hide the interfaces that are associated with the IPv6 family and the interfaces that are not associated with the IPv6 family.

Table 272: TCP Parameters (Continued)

Parameter	Description
	• Deployment type of a Test Agent: • Application—Filter out the Test Agents deployed in a device as an application. • Device—Filter out the Test Agents associated with a device. Based on your requirement, you can select or deselect Application or Device to view or hide the Test Agents deployed in a device as an application or the Test Agents associated with a device. • Connection types between the interface and the router: • Link—Indicates that the status is Up and there is an active connection for the interface. • No Link—Indicates that the status is Down and there is no active connection for the interface. Based on your requirement, you can select or deselect Link or No Link. • Connection statuses of the Test Agent: • Online—The Test Agent is connected to Routing Director. • Offline—The Test Agent is not connected to Routing Director. When the interface supports RPM plugins, the interface contains <i>(On device)</i> at the end of its name indicating the Device Online status.

Table 272: TCP Parameters *(Continued)*

Parameter	Description
	When the interface supports only the normal plugins, the interface will not display <i>(On device)</i> at the end of its name, and the status refers to the Online status. Based on your requirement, you can select or deselect Online or Offline. - Interface types of the Test Agent: - Normal—Network interfaces that are used for testing traffic. - Management—Network interfaces used by the Test Agents to connect to Routing Director. - Unknown—Network interfaces that the Test Agent is unable to identify or that have been removed. Based on your requirement, you can select or deselect Normal, Management, or Unknown. You can enable the Display Device Names toggle button to display the hostnames of the routers that are associated with the Test Agents. You can also enable the Hide unsupported toggle button to hide devices that are not supported by the plugin you have selected. To search for a specific interface or a Test Agent, enter one or more keywords in the search text box. You can search based on the name, description, IP and MAC address of the interfaces or the Test Agents. You can also search based on the device name, device model, and device MAC addresses if searching for a Test Agent associated with the device. You can also hover over an interface name to view the details such as the Test Agent name, bind

Table 272: TCP Parameters *(Continued)*

Parameter	Description
	interface, bind family, and to check if the Test Agent is run on a device or not. You can also remove the interface filters you have selected and reset the filters to the default mode by using the Reset filters button.
Direction	Choose the direction of the test traffic. The data is sent from 'Server to Client' or 'Client to Server'. Client to server—Evaluates the data upload capabilities. Server to client—Evaluates the data download capabilities.
Rate (Mbit/s)	Rate indicates the target data rate between the server and the client Test Agents. ES Rate is the number of errored-seconds occurred when the rate at which the TCP packet reception drops below the defined value. Unit—Megabit per second (Mbit/s) Range—0.01 Mbit/s through 1000000.0 Mbit/s
Advanced	
Server TCP port	Specify the port number used by the server Test Agent. Default value—5000 Range—1 through 65535
Client TCP port	Specify the port number used by the client Test Agent. Range—1 through 65535

Table 272: TCP Parameters *(Continued)*

Parameter	Description
Number of flows	Specify the number of TCP flows. If more than one flows are specified, you should not set the Client port. If set, the plug-in activates an invalid configuration error, and if not set, the Client port gets a randomly assigned port number. Default value—1 Range—1 through 64
DSCP/IPP	Specify the Differentiated Services Code Point (DSCP) or the IP Precedence (IPP) value that is used in the IP packet headers. DSCP is a six-bit binary value in the DS field of the IP header. DSCP value facilitates QoS for traffic management through the Best Effort, Assured Forwarding, Class Selector, and the Expedited Forwarding categories. See RFC 2474 for more information. Default value—0 Range—0 through 63
VLAN priority (PCP)	Specify the value for the VLAN priority code point to be used in VLAN header. Range—0 through 7
Client connect delay (s)	Specify the allowable delay until the client initiates TCP connection towards the server. Unit— Seconds (s) Default value—0 Range—0 through 10
Proxy address	Specify the IP address of the TCP proxy server.

Table 272: TCP Parameters *(Continued)*

Parameter	Description
Proxy port	Specify the port number that the Test Agent uses for TCP proxy server. Range—1 through 65535
Thresholds for error seconds (ES)	
Min rate (Mbit/s)	Specify the minimum expected rate. An errored-second is raised if the TCP data rate drops below the defined value. Unit—Megabit per second (Mbit/s) Minimum—0.001 Mbit/s
Max rate (Mbit/s)	Specify the minimum expected rate. An errored-second is raised if the TCP data rate drops below the defined value. Unit—Megabit per second (Mbit/s) Minimum—0.001 Mbit/s

Table 273: UDP Parameters

Parameter	Description
General	

Table 273: UDP Parameters *(Continued)*

Parameter	Description
Clients*	Select one or more Test Agent interfaces which you want to use as UDP client. To select one or more Test Agent interfaces: 1. Click Clients. The Test Agent/Interfaces choices page appears. 2. Select one or more Test Agent interfaces. For information on filtering Test Agent interfaces, see "Search for Test Agent interfaces by using Filters" on page 1020. 3. Click OK. NOTE: On the Test Agent/Interfaces choices page, you can filter the Test Agent interfaces based on the following criteria: • Interface name or the Test Agent name, or both. Based on your requirement, you can select or deselect Test Agents or Interfaces. • IPv4 address (with or without IPv4 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv4 or No IPv4 to view or hide the interfaces that are associated with the IPv4 family and the interfaces that are not associated with the IPv4 family. • IPv6 address (with or without IPv6 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv6 or No IPv6 to view or hide the interfaces that are associated with the IPv6 family and the interfaces that are not associated with the IPv6 family.

Table 273: UDP Parameters *(Continued)*

Parameter	Description
	- Deployment type of a Test Agent: - Application—Filter out the Test Agents deployed in a device as an application. - Device—Filter out the Test Agents associated with a device. Based on your requirement, you can select or deselect Application or Device to view or hide the Test Agents deployed in a device as an application or the Test Agents associated with a device. - Connection types between the interface and the router: - Link—Indicates that the status is Up and there is an active connection for the interface. - No Link—Indicates that the status is Down and there is no active connection for the interface. Based on your requirement, you can select or deselect Link or No Link. - Connection statuses of the Test Agent: - Online—The Test Agent is connected to Routing Director. - Offline—The Test Agent is not connected to Routing Director. When the interface supports RPM plugins, the interface contains <i>(On device)</i> at the end of its name indicating the Device Online status.

Table 273: UDP Parameters *(Continued)*

Parameter	Description
	When the interface supports only the normal plugins, the interface will not display <i>(On device)</i> at the end of its name, and the status refers to the Online status. Based on your requirement, you can select or deselect Online or Offline. - Interface types of the Test Agent: - Normal—Network interfaces that are used for testing traffic. - Management—Network interfaces used by the Test Agents to connect to Routing Director. - Unknown—Network interfaces that the Test Agent is unable to identify or that have been removed. Based on your requirement, you can select or deselect Normal, Management, or Unknown. You can enable the Display Device Names toggle button to display the hostnames of the routers that are associated with the Test Agents. You can also enable the Hide unsupported toggle button to hide devices that are not supported by the plugin you have selected. To search for a specific interface or a Test Agent, enter one or more keywords in the search text box. You can search based on the name, description, IP and MAC address of the interfaces or the Test Agents. You can also search based on the device name, device model, and device MAC addresses if searching for a Test Agent associated with the device. You can also hover over an interface name to view the details such as the Test Agent name, bind

Table 273: UDP Parameters *(Continued)*

Parameter	Description
	interface, bind family, and to check if the Test Agent is run on a device or not. You can also remove the interface filters you have selected and reset the filters to the default mode by using the Reset filters button.

Table 273: UDP Parameters *(Continued)*

Parameter	Description
Server*	Select one or more Test Agent interfaces which you want to use as UDP server. To select one or more Test Agent interfaces: 1. Click Clients. The Test Agent/Interfaces choices page appears. 2. Select one or more Test Agent interfaces. For information on filtering Test Agent interfaces, see "Search for Test Agent interfaces by using Filters" on page 1024. 3. Click OK. NOTE: On the Test Agent/Interfaces choices page, you can filter the Test Agent interfaces based on the following criteria: • Interface name or the Test Agent name, or both. Based on your requirement, you can select or deselect Test Agents or Interfaces. • IPv4 address (with or without IPv4 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv4 or No IPv4 to view or hide the interfaces that are associated with the IPv4 family and the interfaces that are not associated with the IPv4 family. • IPv6 address (with or without IPv6 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv6 or No IPv6 to view or hide the interfaces that are associated with the IPv6 family and the interfaces that are not associated with the IPv6 family.

Table 273: UDP Parameters *(Continued)*

Parameter	Description
	• Deployment type of a Test Agent: • Application—Filter out the Test Agents deployed in a device as an application. • Device—Filter out the Test Agents associated with a device. Based on your requirement, you can select or deselect Application or Device to view or hide the Test Agents deployed in a device as an application or the Test Agents associated with a device. • Connection types between the interface and the router: • Link—Indicates that the status is Up and there is an active connection for the interface. • No Link—Indicates that the status is Down and there is no active connection for the interface. Based on your requirement, you can select or deselect Link or No Link. • Connection statuses of the Test Agent: • Online—The Test Agent is connected to Routing Director. • Offline—The Test Agent is not connected to Routing Director. When the interface supports RPM plugins, the interface contains <i>(On device)</i> at the end of its name indicating the Device Online status.

Table 273: UDP Parameters *(Continued)*

Parameter	Description
	When the interface supports only the normal plugins, the interface will not display <i>(On device)</i> at the end of its name, and the status refers to the Online status. Based on your requirement, you can select or deselect Online or Offline. - Interface types of the Test Agent: - Normal—Network interfaces that are used for testing traffic. - Management—Network interfaces used by the Test Agents to connect to Routing Director. - Unknown—Network interfaces that the Test Agent is unable to identify or that have been removed. Based on your requirement, you can select or deselect Normal, Management, or Unknown. You can enable the Display Device Names toggle button to display the hostnames of the routers that are associated with the Test Agents. You can also enable the Hide unsupported toggle button to hide devices that are not supported by the plugin you have selected. To search for a specific interface or a Test Agent, enter one or more keywords in the search text box. You can search based on the name, description, IP and MAC address of the interfaces or the Test Agents. You can also search based on the device name, device model, and device MAC addresses if searching for a Test Agent associated with the device. You can also hover over an interface name to view the details such as the Test Agent name, bind

Table 273: UDP Parameters *(Continued)*

Parameter	Description
	interface, bind family, and to check if the Test Agent is run on a device or not. You can also remove the interface filters you have selected and reset the filters to the default mode by using the Reset filters button.
Direction	Choose the direction of the test traffic. The data is sent from 'Server to Client' or 'Client to Server'. Client to server—Evaluates the data upload capabilities. Server to client—Evaluates the data download capabilities.
Rate (Mbit/s)	Specify the target data rate at which the client or server Test Agent transmits the UDP frames. Unit—Megabit per second (Mbit/s) Range—0.01 Mbit/s through 1000000.0 Mbit/s
Advanced	
Server UDP port	Specify the port number of the UDP server to which the client Test Agent form a connection. Default value—5000 Range—1 through 65535
Client TCP port	Specify the port number of the UDP server to which the client Test Agent form a connection. Range—1 through 65535

Table 273: UDP Parameters *(Continued)*

Parameter	Description
Number of flows	Specify the total number of UDP connections established between the client and server Test Agents. If more than one flows are specified, you should not set the Client port. If set, the plug-in activates an invalid configuration error, and if not set, the Client port gets a randomly assigned port number. Default value—1 Range—1 through 64
Don't Fragment flag	Enable the Don't Fragment Flag (DF Flag) to restrict the fragmentation of the packets that exceed the MTU. DF Flag is configured in an IP header. Router drops the packet if fragmentation is needed. Enabling the toggle button may cause performance degradation both in the network and in the sending or receiving Test Agents. By default, the toggle button is disabled.
Ethernet frame size (Bytes)	Specify the size of Layer 2 Ethernet frame for the data flow. Frame Size indicates the total size of the data frame sent from a remote endpoint to the Test agent as a response to a request. Default—1518 Range—64 bytes through 9018 bytes

Table 273: UDP Parameters *(Continued)*

Parameter	Description
DSCP/IPP	Specify the Differentiated Services Code Point (DSCP) or the IP Precedence (IPP) value that is used in the IP packet headers. DSCP is a six-bit binary value in the DS field of the IP header. DSCP value facilitates QoS for traffic management through the Best Effort, Assured Forwarding, Class Selector, and the Expedited Forwarding categories. See RFC 2474 for more information. Default value—0 Range—0 through 63
VLAN priority (PCP)	Specify the value for the VLAN priority code point to be used in VLAN header. Range—0 through 7
Thresholds for error seconds (ES)	
Loss (%)	Specify the loss percentage value. If the loss percentage exceeds the configured value, an errored-second is triggered. Unit—Percentage (%) Range—0.0% through 100.0%
Delay (ms)	Specify the threshold value for the minimum delay measured between the Test Agent client and server. If delay exceeds the defined value, an errored-seconds is raised. Unit—Milliseconds (ms) Minimum—0.0 ms

Table 273: UDP Parameters *(Continued)*

Parameter	Description
Jitter (ms)	Specify the maximum threshold value for delay variance (jitter) in UDP transaction. Delay variation (DV) occurs when different packets take different amount of time to travel from a Test Agent to a remote endpoint. Unit—Milliseconds (ms) Minimum—0.0 ms
Expected DSCP	Specify the expected DSCP that must be in the frame when received by Test Agent. If the received DSCP value does not match the configured value, an errored-second is raised. Range—0 through 64
Thresholds for severely error seconds (SES)	
Loss (%)	Specify the maximum allowable loss percentage value. If the loss percentage exceeds the configured value, an errored-second is raised. Unit—Percentage (%) Range—0.0% through 100.0%
Delay (ms)	Specify the maximum allowable value for delay measurement. If the delay between server and client exceeds the configured value during a one-second interval, a severely-error seconds is indicated. Unit—Milliseconds (ms) Minimum—0.0 ms

Table 273: UDP Parameters *(Continued)*

Parameter	Description
Jitter (ms)	Specify the maximum threshold value for delay variance (jitter). Packets are sent at regular interval of time and if variation is experienced in consecutive packets, the Test Agent generates an errored-second event. Unit—Milliseconds (ms) Minimum—0.0 ms

Table 274: TWAMP/TWAMP Light Parameters

Parameter	Description
General	

Table 274: TWAMP/TWAMP Light Parameters (*Continued*)

Parameter	Description
Senders*	Select one or more Test Agent interfaces on which you want to run the Test. To select one or more Test Agent interfaces: 1. Click the Clients text box. The Test Agent/Interfaces choices page appears. 2. Select one or more Test Agent interfaces. For information on filtering Test Agent interfaces, see "Search for Test Agent interfaces by using Filters" on page 1032. 3. Click OK. NOTE: On the Test Agent/Interfaces choices page, you can filter the Test Agent interfaces based on the following criteria: • Interface name or the Test Agent name, or both. Based on your requirement, you can select or deselect Test Agents or Interfaces. • IPv4 address (with or without IPv4 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv4 or No IPv4 to view or hide the interfaces that are associated with the IPv4 family and the interfaces that are not associated with the IPv4 family. • IPv6 address (with or without IPv6 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv6 or No IPv6 to view or hide the interfaces that are associated with the IPv6 family and the interfaces that are not associated with the IPv6 family.

Table 274: TWAMP/TWAMP Light Parameters (Continued)

Parameter	Description
	• Deployment type of a Test Agent: • Application—Filter out the Test Agents deployed in a device as an application. • Device—Filter out the Test Agents associated with a device. Based on your requirement, you can select or deselect Application or Device to view or hide the Test Agents deployed in a device as an application or the Test Agents associated with a device. • Connection types between the interface and the router: • Link—Indicates that the status is Up and there is an active connection for the interface. • No Link—Indicates that the status is Down and there is no active connection for the interface. Based on your requirement, you can select or deselect Link or No Link. • Connection statuses of the Test Agent: • Online—The Test Agent is connected to Routing Director. • Offline—The Test Agent is not connected to Routing Director. When the interface supports RPM plugins, the interface contains (<i>On device</i>) at the end of its name indicating the Device Online status.

Table 274: TWAMP/TWAMP Light Parameters (*Continued*)

Parameter	Description
	When the interface supports only the normal plugins, the interface will not display (<i>On device</i>) at the end of its name, and the status refers to the Online status. Based on your requirement, you can select or deselect Online or Offline. - Interface types of the Test Agent: - Normal—Network interfaces that are used for testing traffic. - Management—Network interfaces used by the Test Agents to connect to Routing Director. - Unknown—Network interfaces that the Test Agent is unable to identify or that have been removed. Based on your requirement, you can select or deselect Normal, Management, or Unknown. You can enable the Display Device Names toggle button to display the hostnames of the routers that are associated with the Test Agents. You can also enable the Hide unsupported toggle button to hide devices that are not supported by the plugin you have selected. To search for a specific interface or a Test Agent, enter one or more keywords in the search text box. You can search based on the name, description, IP and MAC address of the interfaces or the Test Agents. You can also search based on the device name, device model, and device MAC addresses if searching for a Test Agent associated with the device. You can also hover over an interface name to view the details such as the Test Agent name, bind

Table 274: TWAMP/TWAMP Light Parameters *(Continued)*

Parameter	Description
	interface, bind family, and to check if the Test Agent is run on a device or not. You can also remove the interface filters you have selected and reset the filters to the default mode by using the Reset filters button. The TWAMP/TWAMP Light plugin cannot be run on Test Agents associated with devices. To run TWAMP/TWAMP Light plugin on devices, use the RPM TWAMP/TWAMP Light plugin.

Table 274: TWAMP/TWAMP Light Parameters (Continued)

Parameter	Description
Reflectors*	Specify the reflector address. A Test Agent application can run a Reflector plugin whereas a Test Agent that is associated with a device needs to be configured to run Reflector plugin. When you click the text box, the Select reflectors page appears where you can add reflectors. On this page: • Reflector hostname—Specify the hostname for the reflector. The reflector receives the TWAMP request packets and reflect the packet back to the sender. Maximum length—64 characters • Test session port—Specify the destination port for the Test session. Range—0 through 65535 • Control session port—Specify the port for the control session. If set to zero, TWAMP Light is used. Default value—0 Range—0 through 65535 • Source port for Test session port—Specify the port for the control session. If set to zero, TWAMP Light is used. Specify the source port value for the Test session. Default value—0 Range—0 through 65535 • Name—This text box is automatically populated based on the data you specified in the Reflector hostname text box.

Table 274: TWAMP/TWAMP Light Parameters (*Continued*)

Parameter	Description
Rate	Specify the rate at which the TWAMP sender will send test traffic. The rate is calculated as the size of the request packet sent divided by the total request time. Each Ethernet packet contains one frame. Unit—Megabits per seconds (Mbit/s) Range—0.0 Mbit/s through 10000.0 Mbit/s
Time sync	Select the time synchronization option. Select Yes, if the clocks of Test Agent and the reflector are synchronized via Network Time Protocol (NTP) using external NTP servers. Select No to disable the time synchronization option. Select Inband to use the in-band synchronization algorithm.
Use hardware timestamping	Enable hardware timestamping, SO_TIMESTAMPING parameter, on TWAMP socket), if you want to use the network interface card (NIC) of Test Agents for delay and jitter measurements. If Test Agent NIC does not support it, an error message is displayed, and the measurement does not produce any delay metrics. By default, the toggle button is disabled.
Thresholds for error seconds (ES)	

Table 274: TWAMP/TWAMP Light Parameters (*Continued*)

Parameter	Description
Loss%	Specify the loss percentage value. If the loss percentage exceeds the configured value, an errored-second is raised. The Loss percentage is calculated by comparing the total number of packets that were lost with the total number of packets that were sent from the Test Agent. Unit—Percentage (%) Default value—0.0 % Range—0.0 % through 100.00 %
Delay	Specify the maximum tolerated value for delay. If there is time synchronization, the threshold value applies only to one-way delay. If there is no time synchronization, the threshold applies to round-trip delay. If delay exceeds the defined value, an errored-second is raised. Unit—Milliseconds (ms) Range—0.001 ms through 1000.0 ms
Delay Variation	Specify the maximum tolerated value for delay variance. If there is time synchronization, it applies to one-way delay. If there is no time synchronization, it applies to round-trip delay. If delay exceeds the defined value, an errored-second is raised. Unit—Milliseconds (ms) Range—0.001 ms through 1000.0 ms

Table 274: TWAMP/TWAMP Light Parameters (*Continued*)

Parameter	Description
Expected DSCP Value	Specify the expected DSCP or IP Precedence value on the IP packets received from the reflector device. If the received DSCP value does not match the configured value, an errored-second will be raised. Range—0 through 63
Thresholds for severely error seconds (SES)	
Loss	Specify the loss percentage value. If the loss percentage exceeds the configured value during a one-second interval, a severely errored-second is raised. The Loss percentage is calculated by comparing the total number of packets that were lost with the total number of packets that were sent from the Test Agent. Unit—Percentage (%) Minimum value—0.0 %
Delay	Specify the maximum acceptable delay. If there is time synchronization, it applies to one-way delay. If there is no time synchronization, it applies to round-trip delay. If delay exceeds the defined value, a severely errored-second is raised. Unit—Milliseconds (ms) Minimum value—0.001 ms

Table 274: TWAMP/TWAMP Light Parameters (*Continued*)

Parameter	Description
Delay variation	Specify the maximum tolerated value for delay variance (jitter). If there is time synchronization, it applies to one-way delay. If there is no time synchronization, it applies to round-trip delay. If delay exceeds the defined value, a severely errored-second is raised. Unit—Milliseconds (ms) Minimum value—0.001 ms
Advanced	
Frame Size	Specify the size of Layer 2 Ethernet frame size for the TWAMP test traffic. Frame Size indicates the total size of the data frame sent from the Test Agent to the remote endpoint. The size also includes the header size. Unit—Bytes Default value—1518 Range—87 through 9018
DSCP	Specify the Differentiated Services Code Point (DSCP) or IP Precedence value that is used in the IP packet headers. Default value—0 Range—0 through 63
Use random padding	Enable this to use random numbers or zeroes as padding in a TWAMP packet. By default, the toggle button is enabled.

Table 274: TWAMP/TWAMP Light Parameters *(Continued)*

Parameter	Description
Socket priority	Specify the value of SO_PRIORITY parameter of the TWAMP socket. Default value—0 Range—0 through 7
Socket send buffer size	Specify the value of SO_PRIORITY parameter of the TWAMP socket. Unit—Bytes Range—2048 through 10000000 bytes
Socket receive buffer size	Specify the value of SO_PRIORITY parameter of the TWAMP socket. Unit—Bytes Range—2048 through 10000000 bytes
Don't fragment flag	Enable the Don't Fragment Flag (DF Flag) to restrict the fragmentation of the packets that exceed the MTU. DF Flag is configured in an IP header. Router drops the packet if fragmentation is needed. Enabling this option may cause performance degradation both in the network and in the sending or receiving Test Agents. By default, the toggle button is enabled.

Table 274: TWAMP/TWAMP Light Parameters (*Continued*)

Parameter	Description
UAS period length	Specify the minimum value for the consecutive severely errored-seconds (SES) that causes a period of unavailability. The Unavailable Seconds (UAS) metric determines the number of seconds at which the service can be considered to be unavailable. Unit—Seconds (s) Default value—10 s Range—0 s through 300 s
Accept UDP checksum zero for IPv6	Enable Accept UDP Checksum Zero for IPv6 to accept the UDP Checksum as Zero for IPv6 in TWAMP Reflector packets. By default, the toggle button is enabled.
Percentiles	
First delay percentile	Specify the first delay percentile of the value of the TWAMP request packet. If the delay exceeds the configured value, the packet is included in the defined first delay percentile slot. Unit—Percentage (%) Range—0 % through 1 %
Second delay percentile	Specify the second delay percentile of the value of the TWAMP request packet. If the delay exceeds the configured value, the packet is included in the defined second delay percentile slot. Unit—Percentage (%) Range—0 % through 1 %

Table 274: TWAMP/TWAMP Light Parameters (*Continued*)

Parameter	Description
Threshold for first delay percentile	Specify the threshold for triggering an errored second based on the first delay percentile. Unit—Milliseconds (ms). Range—0.001 ms through 1000.00 ms
Threshold for second delay percentile	Specify the threshold for triggering an errored second based on the second delay percentile. Unit—Milliseconds (ms). Range—0.001 ms through 1000.00 ms
SES threshold for first delay percentile	Specify the threshold for triggering a severely errored second based on the first delay percentile. Unit—Milliseconds (ms). Range—0.001 ms through 1000.00 ms
SES threshold for second delay percentile	Specify the threshold for triggering a severely errored second based on the second delay percentile. Unit—Milliseconds (ms). Range—0.001 ms through 1000.00 ms
Periodic Streams	
Active period duration	Specify the time duration of each cycle during which ethernet frames are sent. Active period is followed by a silent period during which no ethernet frames are sent. Unit—Milliseconds (ms) Range—1 ms to 3600000 ms

Table 274: TWAMP/TWAMP Light Parameters (*Continued*)

Parameter	Description
Active cycle	Specify the time duration of the cycle starting with an active period and ending with a silent period. The Active cycle duration must be at least equal to the active period duration. Unit—Milliseconds (ms) Range—1 ms to 604800 ms
Report metrics during inactive period	Enable the toggle button to report metrics related to inactive periods of a periodic Test. By default, the toggle button is enabled.

Table 275: TWAMP Reflector Parameters

Parameter	Description
General	

Table 275: TWAMP Reflector Parameters *(Continued)*

Parameter	Description
Clients*	Select one or more Test Agent interfaces on which you want to run the Test. To select one or more Test Agent interfaces: 1. Click the Clients text box. The Test Agent/Interfaces choices page appears. 2. Select one or more Test Agent interfaces. For information on filtering Test Agent interfaces, see "Search for Test Agent interfaces by using Filters" on page 1045. 3. Click OK. NOTE: On the Test Agent/Interfaces choices page, you can filter the Test Agent interfaces based on the following criteria: • Interface name or the Test Agent name, or both. Based on your requirement, you can select or deselect Test Agents or Interfaces. • IPv4 address (with or without IPv4 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv4 or No IPv4 to view or hide the interfaces that are associated with the IPv4 family and the interfaces that are not associated with the IPv4 family. • IPv6 address (with or without IPv6 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv6 or No IPv6 to view or hide the interfaces that are associated with the IPv6 family and the interfaces that are not associated with the IPv6 family.

Table 275: TWAMP Reflector Parameters (Continued)

Parameter	Description
	- Deployment type of a Test Agent: - Application—Filter out the Test Agents deployed in a device as an application. - Device—Filter out the Test Agents associated with a device. Based on your requirement, you can select or deselect Application or Device to view or hide the Test Agents deployed in a device as an application or the Test Agents associated with a device. - Connection types between the interface and the router: - Link—Indicates that the status is Up and there is an active connection for the interface. - No Link—Indicates that the status is Down and there is no active connection for the interface. Based on your requirement, you can select or deselect Link or No Link. - Connection statuses of the Test Agent: - Online—The Test Agent is connected to Routing Director. - Offline—The Test Agent is not connected to Routing Director. When the interface supports RPM plugins, the interface contains <i>(On device)</i> at the end of its name indicating the Device Online status.

Table 275: TWAMP Reflector Parameters (Continued)

Parameter	Description
	When the interface supports only the normal plugins, the interface will not display (<i>On device</i>) at the end of its name, and the status refers to the Online status. Based on your requirement, you can select or deselect Online or Offline. - Interface types of the Test Agent: - Normal—Network interfaces that are used for testing traffic. - Management—Network interfaces used by the Test Agents to connect to Routing Director. - Unknown—Network interfaces that the Test Agent is unable to identify or that have been removed. Based on your requirement, you can select or deselect Normal, Management, or Unknown. You can enable the Display Device Names toggle button to display the hostnames of the routers that are associated with the Test Agents. You can also enable the Hide unsupported toggle button to hide devices that are not supported by the plugin you have selected. To search for a specific interface or a Test Agent, enter one or more keywords in the search text box. You can search based on the name, description, IP and MAC address of the interfaces or the Test Agents. You can also search based on the device name, device model, and device MAC addresses if searching for a Test Agent associated with the device. You can also hover over an interface name to view the details such as the Test Agent name, bind

Table 275: TWAMP Reflector Parameters *(Continued)*

Parameter	Description
	interface, bind family, and to check if the Test Agent is run on a device or not. You can also remove the interface filters you have selected and reset the filters to the default mode by using the Reset filters button.
Test session port	Specify the destination port on which the TWAMP reflector listens. Default value—7000 Range—1 through 65535
Rate threshold for ES	Specify the rate threshold for the combined rate for all sessions. If rate is below the threshold value, an errored-second is raised. Unit—Megabits per seconds (Mbit/s) Range—0.001 (Mbit/s) through 10000 (Mbit/s)
Standalone mode	Enable Standalone mode to push the metrics data to Routing Director. By default, the toggle button is disabled.

RPM Plug-ins

Table 276: RPM HTTP Parameters

Parameter	Description
General	

Table 276: RPM HTTP Parameters *(Continued)*

Parameter	Description
Clients*	Select one or more interfaces of the network devices on which you want to run a Test. To select one or more network devices as client values: 1. Click the Clients text box. The Test Agent/Interfaces choices page appears. 2. Select one or more Test Agent interfaces. For information on filtering Test Agent interfaces, see "Search for Test Agent interfaces by using Filters" on page 1049. 3. Click OK. NOTE: On the Test Agent/Interfaces choices page, you can filter the Test Agent interfaces based on the following criteria: • Interface name or the Test Agent name, or both. Based on your requirement, you can select or deselect Test Agents or Interfaces. NOTE: The Test Agent filter is not functional as the RPM plugins directly interact with the devices. • IPv4 address (with or without IPv4 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv4 or No IPv4 to view or hide the interfaces that are associated with the IPv4 family and the interfaces that are not associated with the IPv4 family. • IPv6 address (with or without IPv6 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv6 or No IPv6 to view or hide the

Table 276: RPM HTTP Parameters *(Continued)*

Parameter	Description
	interfaces that are associated with the IPv6 family and the interfaces that are not associated with the IPv6 family. - Deployment type of a Test Agent: - Application—Filter out the Test Agents deployed in a device as an application. - Device—Filter out the Test Agents associated with a device. Based on your requirement, you can select or deselect Application or Device to view or hide the Test Agents deployed in a device as an application or the Test Agents associated with a device. - Connection types between the interface and the router: - Link—Indicates that the status is Up and there is an active connection for the interface. - No Link—Indicates that the status is Down and there is no active connection for the interface. Based on your requirement, you can select or deselect Link or No Link. - Connection statuses of the Test Agent: - Online—The Test Agent is connected to Routing Director. - Offline—The Test Agent is not connected to Routing Director.

Table 276: RPM HTTP Parameters *(Continued)*

Parameter	Description
	When the interface supports RPM plugins, the interface contains <i>(On device)</i> at the end of its name indicating the Device Online status. When the interface supports only the normal plugins, the interface will not display <i>(On device)</i> at the end of its name, and the status refers to the Online status. Based on your requirement, you can select or deselect Online or Offline. • Interface types of the Test Agent: • Normal—Network interfaces that are used for testing traffic. • Management—Network interfaces used by the Test Agents to connect to Routing Director. • Unknown—Network interfaces that the Test Agent is unable to identify or that have been removed. Based on your requirement, you can select or deselect Normal, Management, or Unknown. You can enable the Display Device Names toggle button to display the hostnames of the routers that are associated with the Test Agents. You can also enable the Hide unsupported toggle button to hide devices that are not supported by the plugin you have selected. To search for a specific interface or a Test Agent, enter one or more keywords in the search text box. You can search based on the name, description, IP and MAC address of the interfaces or the Test Agents. You can also search based on the device name, device model, and device MAC addresses if

Table 276: RPM HTTP Parameters *(Continued)*

Parameter	Description
	searching for a Test Agent associated with the device. You can also hover over an interface name to view the details such as the Test Agent name, bind interface, bind family, and to check if the Test Agent is run on a device or not. You can also remove the interface filters you have selected and reset the filters to the default mode by using the Reset filters button.
URLs*	Specify the URL or host towards which the router sends HTTP Get requests. Maximum length—255 characters
Time between requests	Specify the time taken between successive HTTP requests initiated by the router. Unit—Seconds (s) Default value—10 s Range—1 s to 255 s
Advanced	
Collection Interval	Specify the collection interval. Collection interval is the frequency at which the results are collected from the device. Collection interval must be larger than the specified Time between requests. Unit—Seconds (s) Default value—15 s Range—5 s through 300 s

Table 276: RPM HTTP Parameters *(Continued)*

Parameter	Description
Device response timeout	Specify the device response timeout value. Device response timeout determines the time to wait for a response from the router. If the specified time exceeds, the connection to the device is considered to be failed. Device response timeout must be larger than the specified Collection Interval. Unit—Seconds (s) Default value—200 s Range—30 s through 300 s
Routing instance	Specify the name of routing instance on the router used by the probes.
TTL	Specify the number of times the packets hop before a device discards the packet. Time to live (TTL) value indicates lifespan of a request packet. TTL prevents infinite loop in a network when the packet fails to reach the destination. Default value—64. Range—1 through 254.
Metadata get	Enable Metadata get for the Test to perform HTTP Get requests for metadata to a target URL. By default, the toggle button is disabled.
IPv6 local link	Specify the router link-local logical interface name for the egress interface to be used. This applies only for Test sessions with IPv6 address as the target address. Maximum length—64 characters

Table 276: RPM HTTP Parameters *(Continued)*

Parameter	Description
Hardware timestamp	Specify the router link-local logical interface name for the egress interface to be used. This applies only for Test sessions with IPv6 address as the target address. By default, the toggle button is disabled.

Table 277: RPM PING Parameters

Parameter	Description
General	

Table 277: RPM PING Parameters (Continued)

Parameter	Description
Clients*	Select one or more interfaces of the network devices on which you want to run a Test. To select one or more network devices as client values: 1. Click the Clients text box. The Test Agent/Interfaces choices page appears. 2. Select one or more Test Agent interfaces. For information on filtering Test Agent interfaces, see "Search for Test Agent interfaces by using Filters" on page 1055. 3. Click OK. NOTE: On the Test Agent/Interfaces choices page, you can filter the Test Agent interfaces based on the following criteria: • Interface name or the Test Agent name, or both. Based on your requirement, you can select or deselect Test Agents or Interfaces. NOTE: The Test Agent filter is not functional as the RPM plugins directly interact with the devices. • IPv4 address (with or without IPv4 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv4 or No IPv4 to view or hide the interfaces that are associated with the IPv4 family and the interfaces that are not associated with the IPv4 family. • IPv6 address (with or without IPv6 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv6 or No IPv6 to view or hide the

Table 277: RPM PING Parameters *(Continued)*

Parameter	Description
	interfaces that are associated with the IPv6 family and the interfaces that are not associated with the IPv6 family. • Deployment type of a Test Agent: • Application—Filter out the Test Agents deployed in a device as an application. • Device—Filter out the Test Agents associated with a device. Based on your requirement, you can select or deselect Application or Device to view or hide the Test Agents deployed in a device as an application or the Test Agents associated with a device. • Connection types between the interface and the router: • Link—Indicates that the status is Up and there is an active connection for the interface. • No Link—Indicates that the status is Down and there is no active connection for the interface. Based on your requirement, you can select or deselect Link or No Link. • Connection statuses of the Test Agent: • Online—The Test Agent is connected to Routing Director. • Offline—The Test Agent is not connected to Routing Director.

Table 277: RPM PING Parameters *(Continued)*

Parameter	Description
	When the interface supports RPM plugins, the interface contains <i>(On device)</i> at the end of its name indicating the Device Online status. When the interface supports only the normal plugins, the interface will not display <i>(On device)</i> at the end of its name, and the status refers to the Online status. Based on your requirement, you can select or deselect Online or Offline. - Interface types of the Test Agent: - Normal—Network interfaces that are used for testing traffic. - Management—Network interfaces used by the Test Agents to connect to Routing Director. - Unknown—Network interfaces that the Test Agent is unable to identify or that have been removed. Based on your requirement, you can select or deselect Normal, Management, or Unknown. You can enable the Display Device Names toggle button to display the hostnames of the routers that are associated with the Test Agents. You can also enable the Hide unsupported toggle button to hide devices that are not supported by the plugin you have selected. To search for a specific interface or a Test Agent, enter one or more keywords in the search text box. You can search based on the name, description, IP and MAC address of the interfaces or the Test Agents. You can also search based on the device name, device model, and device MAC addresses if

Table 277: RPM PING Parameters *(Continued)*

Parameter	Description
	searching for a Test Agent associated with the device. You can also hover over an interface name to view the details such as the Test Agent name, bind interface, bind family, and to check if the Test Agent is run on a device or not. You can also remove the interface filters you have selected and reset the filters to the default mode by using the Reset filters button.
Hosts*	Specify the target hostname or IP address. When you click the text box, the Select Host page appears where you can enter hostnames. To add more than one Hosts, click + Add Host and specify the following: • Host—The hostname or the IP address of the remote endpoint. Maximum length—255 characters. • Name—The text box is automatically populated based on the data you specified in the Host text box.
Time between requests	Specify the time taken between successive ping requests initiated by the router. Unit—Seconds (s) Default value—10 s Range—1 s through 255 s
Advanced	

Table 277: RPM PING Parameters *(Continued)*

Parameter	Description
Collection interval	Specify the collection interval. Collection interval is the frequency at which the results are collected from the device. Collection interval must be larger than the specified Time between requests. Unit—Seconds (s) Default value—15 s Range—5 s through 300 s
Device response timeout	Specify the device response timeout value. Device response timeout determines the time to wait for a response from the router. If the specified time is exceeded, the connection to the device is considered failed. Device response timeout must be larger than the specified Collection Interval. Unit—Seconds (s) Default value—200 s Range—30 s through 300 s
Routing instance	Specify the name of routing instance on the router used by the probes.
Data size	Specify the size of the data portion of the request packet sent from the router. Default value—0 Range—0 through 65400
Data fill	Specify the contents of the data portion of the request packets sent from the router. The value should be in hexadecimal format.

Table 277: RPM PING Parameters *(Continued)*

Parameter	Description
TTL	Specify the number of times the packets hop before a device discards the packet. Time to live (TTL) value indicates lifespan of a request packet. TTL prevents infinite loop in a network when the packet fails to reach the destination. For example, if you have entered the TTL value as 64, every time the packet passes a device, the value is reduced by one until the packet reaches the remote server. If the entered value reaches zero before reaching the remote endpoint, the packet is discarded. Default value—64 Range—1 through 254
DSCP code points	Specify the Differentiated Services Code Point (DSCP) value that is used in the IP packet headers. Maximum length—64 characters
Hardware timestamp	Enable hardware timestamping of RPM probe messages in the Packet Forwarding Engine host processor on the router. In case of ping plugin, you can enable timestamping of RPM messages in the Packet Forwarding Engine host processor. By default, the toggle button is disabled.
Ping timestamp	Enable ping timestamp to perform ICMP timestamp requests instead of ICMP echo requests. By default, the toggle button is disabled.
One way hardware timestamp	Enable the one-way hardware timestamping for one-way measurements (delay and jitter) By default, the toggle button is disabled.

Table 277: RPM PING Parameters *(Continued)*

Parameter	Description
IPv6 local link	Specify the router link-local logical interface name for the egress interface to be used. This applies only for Test sessions with IPv6 address as the target address. Maximum length—64 characters

Table 278: RPM TCP Parameters

Parameter	Description
General	

Table 278: RPM TCP Parameters (Continued)

Parameter	Description
Clients*	Select one or more interfaces of the network devices on which you want to run a Test. To select one or more network devices as client values: 1. Click the Clients text box. The Test Agent/Interfaces choices page appears. 2. Select one or more Test Agent interfaces. For information on filtering Test Agent interfaces, see "Search for Test Agent interfaces by using Filters" on page 1062. 3. Click OK. NOTE: On the Test Agent/Interfaces choices page, you can filter the Test Agent interfaces based on the following criteria: • Interface name or the Test Agent name, or both. Based on your requirement, you can select or deselect Test Agents or Interfaces. NOTE: The Test Agent filter is not functional as the RPM plugins directly interact with the devices. • IPv4 address (with or without IPv4 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv4 or No IPv4 to view or hide the interfaces that are associated with the IPv4 family and the interfaces that are not associated with the IPv4 family. • IPv6 address (with or without IPv6 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv6 or No IPv6 to view or hide the

Table 278: RPM TCP Parameters *(Continued)*

Parameter	Description
	interfaces that are associated with the IPv6 family and the interfaces that are not associated with the IPv6 family. • Deployment type of a Test Agent: • Application—Filter out the Test Agents deployed in a device as an application. • Device—Filter out the Test Agents associated with a device. Based on your requirement, you can select or deselect Application or Device to view or hide the Test Agents deployed in a device as an application or the Test Agents associated with a device. • Connection types between the interface and the router: • Link—Indicates that the status is Up and there is an active connection for the interface. • No Link—Indicates that the status is Down and there is no active connection for the interface. Based on your requirement, you can select or deselect Link or No Link. • Connection statuses of the Test Agent: • Online—The Test Agent is connected to Routing Director. • Offline—The Test Agent is not connected to Routing Director.

Table 278: RPM TCP Parameters (Continued)

Parameter	Description
	When the interface supports RPM plugins, the interface contains <i>(On device)</i> at the end of its name indicating the Device Online status. When the interface supports only the normal plugins, the interface will not display <i>(On device)</i> at the end of its name, and the status refers to the Online status. Based on your requirement, you can select or deselect Online or Offline. - Interface types of the Test Agent: - Normal—Network interfaces that are used for testing traffic. - Management—Network interfaces used by the Test Agents to connect to Routing Director. - Unknown—Network interfaces that the Test Agent is unable to identify or that have been removed. Based on your requirement, you can select or deselect Normal, Management, or Unknown. You can enable the Display Device Names toggle button to display the hostnames of the routers that are associated with the Test Agents. You can also enable the Hide unsupported toggle button to hide devices that are not supported by the plugin you have selected. To search for a specific interface or a Test Agent, enter one or more keywords in the search text box. You can search based on the name, description, IP and MAC address of the interfaces or the Test Agents. You can also search based on the device name, device model, and device MAC addresses if

Table 278: RPM TCP Parameters (Continued)

Parameter	Description
	searching for a Test Agent associated with the device. You can also hover over an interface name to view the details such as the Test Agent name, bind interface, bind family, and to check if the Test Agent is run on a device or not. You can also remove the interface filters you have selected and reset the filters to the default mode by using the Reset filters button.
Servers*	Specify the hostname or the IP address of the remote endpoint. When you click the text box, the Select server page appears where you can enter details. To add more than one servers, click + Add Server and specify the following: - Remote IP or hostname—The hostname or the IP address of the remote endpoint. Maximum length—255 characters - Remote TCP port—Specify the port number of TCP. Default value—7 Range—7 through 65535
Time between requests	Specify the time taken between successive TCP requests initiated by the router. Unit—Seconds (s) Default value—10 s Range—1 s through 255 s
Advanced	

Table 278: RPM TCP Parameters *(Continued)*

Parameter	Description
Collection interval	Specify the collection interval. Collection interval is the frequency at which the results are collected from the device. Collection interval must be larger than the specified Time between requests. Unit—Seconds (s) Default value—15 s Range—5 s through 300 s
Device response timeout	Specify the device response timeout value. Device response timeout determines the time to wait for a response from the router. If the specified time is exceeded, the connection to the device is considered failed. Device response timeout must be larger than the specified Collection Interval. Unit—Seconds (s) Default value—200 s Range—30 s through 300 s
Routing instance	Specify the name of routing instance on the router used by the request packets. Maximum value—64
Data size	Specify the size of the data portion of the request packets sent from the router. Default value—0 Range—0 through 65400
Data fill	Specify the contents of the data portion of the request packets sent from the router. The value should be in hexadecimal format.

Table 278: RPM TCP Parameters *(Continued)*

Parameter	Description
TTL	Specify the number of times the packets hop before a device discards the packet. Time to live (TTL) value indicates lifespan of a request packet. TTL prevents infinite loop in a network when the packet fails to reach the destination. For example, if you have entered the TTL value as 64, every time the packet passes a device, the value is reduced by one until the packet reaches the remote server. If the entered value reaches zero before reaching the remote endpoint, the packet is discarded. Default value—64 Range—1 through 254
DSCP code points	Specify the Differentiated Services Code Point (DSCP) value that is used in the IP packet headers. Maximum length—64 characters
IPv6 Local Link	Specify the router link-local logical interface name for the egress interface to be used. This applies only for Test sessions with IPv6 address as the target address. Maximum length—64 characters

Table 279: RPM TWAMP Parameters

Parameter	Description
General	

Table 279: RPM TWAMP Parameters *(Continued)*

Parameter	Description
Clients*	Select one or more interfaces of the network devices on which you want to run a Test. To select one or more network devices as client values: 1. Click the Clients text box. The Test Agent/Interfaces choices page appears. 2. Select one or more Test Agent interfaces. For information on filtering Test Agent interfaces, see "Search for Test Agent interfaces by using Filters" on page 1068. 3. Click OK. NOTE: On the Test Agent/Interfaces choices page, you can filter the Test Agent interfaces based on the following criteria: • Interface name or the Test Agent name, or both. Based on your requirement, you can select or deselect Test Agents or Interfaces. NOTE: The Test Agent filter is not functional as the RPM plugins directly interact with the devices. • IPv4 address (with or without IPv4 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv4 or No IPv4 to view or hide the interfaces that are associated with the IPv4 family and the interfaces that are not associated with the IPv4 family. • IPv6 address (with or without IPv6 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv6 or No IPv6 to view or hide the

Table 279: RPM TWAMP Parameters *(Continued)*

Parameter	Description
	interfaces that are associated with the IPv6 family and the interfaces that are not associated with the IPv6 family. • Deployment type of a Test Agent: • Application—Filter out the Test Agents deployed in a device as an application. • Device—Filter out the Test Agents associated with a device. Based on your requirement, you can select or deselect Application or Device to view or hide the Test Agents deployed in a device as an application or the Test Agents associated with a device. • Connection types between the interface and the router: • Link—Indicates that the status is Up and there is an active connection for the interface. • No Link—Indicates that the status is Down and there is no active connection for the interface. Based on your requirement, you can select or deselect Link or No Link. • Connection statuses of the Test Agent: • Online—The Test Agent is connected to Routing Director. • Offline—The Test Agent is not connected to Routing Director.

Table 279: RPM TWAMP Parameters *(Continued)*

Parameter	Description
	When the interface supports RPM plugins, the interface contains <i>(On device)</i> at the end of its name indicating the Device Online status. When the interface supports only the normal plugins, the interface will not display <i>(On device)</i> at the end of its name, and the status refers to the Online status. Based on your requirement, you can select or deselect Online or Offline. - Interface types of the Test Agent: - Normal—Network interfaces that are used for testing traffic. - Management—Network interfaces used by the Test Agents to connect to Routing Director. - Unknown—Network interfaces that the Test Agent is unable to identify or that have been removed. Based on your requirement, you can select or deselect Normal, Management, or Unknown. You can enable the Display Device Names toggle button to display the hostnames of the routers that are associated with the Test Agents. You can also enable the Hide unsupported toggle button to hide devices that are not supported by the plugin you have selected. To search for a specific interface or a Test Agent, enter one or more keywords in the search text box. You can search based on the name, description, IP and MAC address of the interfaces or the Test Agents. You can also search based on the device name, device model, and device MAC addresses if

Table 279: RPM TWAMP Parameters *(Continued)*

Parameter	Description
	searching for a Test Agent associated with the device. You can also hover over an interface name to view the details such as the Test Agent name, bind interface, bind family, and to check if the Test Agent is run on a device or not. You can also remove the interface filters you have selected and reset the filters to the default mode by using the Reset filters button.

Table 279: RPM TWAMP Parameters (Continued)

Parameter	Description
Hosts*	Specify IPv4 or IPv6 Address of TWAMP reflector. When you click the text box, the Select reflectors page appears where you can add reflectors. On this page: • Reflector hostname—Specify the hostname for the reflector. Maximum length—64 characters • Test session port—Specify IPv4 or IPv6 Address of TWAMP reflector. Range—0 through 65535 • Control session port—Specify the port for the control session. If the port value is zero, it is for TWAMP Light, and if port value is in the range of 862-65535, it is for managed TWAMP. Default value—0 Range—0 through 65535 • Source port for Test session port—Specify the source port value for the Test session. Default value—0 Range—0 through 65535 • Name—The text box is automatically populated based on the data you specified in the Reflector hostname text box.
Time between requests*	Specify the time taken between successive TWAMP requests initiated by the router. Unit—Seconds (s) Default value—10 s Range—1 s through 255 s

Table 279: RPM TWAMP Parameters *(Continued)*

Parameter	Description
Advanced	
Collection interval*	Specify the collection interval. Collection interval is the frequency at which the results are collected from the device. Collection interval must be larger than the specified Time between requests. Unit—Seconds (s) Default value—15 s Range—5 s through 300 s
Device response timeout	Specify the device response timeout value. Device response timeout determines the time to wait for a response from the router. If the specified time is exceeded, the connection to the device is considered failed. Device response timeout must be larger than the specified Collection Interval. Unit—Seconds (s) Default value—200 s Range—30 s through 300 s
Routing instance	Specify the name of routing instance on the router used by the request packets. Maximum value—64
Zero fill	Enable Zero Fill to populate the content for the request packet with zeros.

Table 279: RPM TWAMP Parameters *(Continued)*

Parameter	Description
Data size	Specify the size of the data portion of the request packets sent from the router. Default value—60. Range—60 to 1400.
TTL	Specify the number of times the packets hop before a device discards the packet. Time to live (TTL) value indicates lifespan of a request packet. TTL prevents infinite loop in a network when the packet fails to reach the destination. For example, if you have entered the TTL value as 64, every time the packet passes a device, the value is reduced by one until the packet reaches the remote server. If the entered value reaches zero before reaching the remote endpoint, the packet is discarded. Default value—64 Range—1 to 254
DSCP Code Points	Specify the Differentiated Services Code Point (DSCP) value that is used in the IP packet headers. Maximum length—64 characters
IPv6 Local Link	Specify the router link-local logical interface name for the egress interface to be used. This applies only for Test sessions with IPv6 address as the target address. Maximum length—64 characters
PFE timestamping	Enable PFE timestamping to perform timestamping on Packet Forward Engine host. By default, the toggle button is disabled.

Table 280: RPM UDP Parameters

Parameter	Description
General	

Table 280: RPM UDP Parameters *(Continued)*

Parameter	Description
Clients*	Select one or more interfaces of the network devices on which you want to run a Test. To select one or more network devices as client values: 1. Click the Clients text box. The Test Agent/Interfaces choices page appears. 2. Select one or more Test Agent interfaces. For information on filtering Test Agent interfaces, see "Search for Test Agent interfaces by using Filters" on page 1076. 3. Click OK. NOTE: On the Test Agent/Interfaces choices page, you can filter the Test Agent interfaces based on the following criteria: • Interface name or the Test Agent name, or both. Based on your requirement, you can select or deselect Test Agents or Interfaces. NOTE: The Test Agent filter is not functional as the RPM plugins directly interact with the devices. • IPv4 address (with or without IPv4 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv4 or No IPv4 to view or hide the interfaces that are associated with the IPv4 family and the interfaces that are not associated with the IPv4 family. • IPv6 address (with or without IPv6 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv6 or No IPv6 to view or hide the

Table 280: RPM UDP Parameters *(Continued)*

Parameter	Description
	interfaces that are associated with the IPv6 family and the interfaces that are not associated with the IPv6 family. • Deployment type of a Test Agent: • Application—Filter out the Test Agents deployed in a device as an application. • Device—Filter out the Test Agents associated with a device. Based on your requirement, you can select or deselect Application or Device to view or hide the Test Agents deployed in a device as an application or the Test Agents associated with a device. • Connection types between the interface and the router: • Link—Indicates that the status is Up and there is an active connection for the interface. • No Link—Indicates that the status is Down and there is no active connection for the interface. Based on your requirement, you can select or deselect Link or No Link. • Connection statuses of the Test Agent: • Online—The Test Agent is connected to Routing Director. • Offline—The Test Agent is not connected to Routing Director.

Table 280: RPM UDP Parameters *(Continued)*

Parameter	Description
	When the interface supports RPM plugins, the interface contains <i>(On device)</i> at the end of its name indicating the Device Online status. When the interface supports only the normal plugins, the interface will not display <i>(On device)</i> at the end of its name, and the status refers to the Online status. Based on your requirement, you can select or deselect Online or Offline. - Interface types of the Test Agent: - Normal—Network interfaces that are used for testing traffic. - Management—Network interfaces used by the Test Agents to connect to Routing Director. - Unknown—Network interfaces that the Test Agent is unable to identify or that have been removed. Based on your requirement, you can select or deselect Normal, Management, or Unknown. You can enable the Display Device Names toggle button to display the hostnames of the routers that are associated with the Test Agents. You can also enable the Hide unsupported toggle button to hide devices that are not supported by the plugin you have selected. To search for a specific interface or a Test Agent, enter one or more keywords in the search text box. You can search based on the name, description, IP and MAC address of the interfaces or the Test Agents. You can also search based on the device name, device model, and device MAC addresses if

Table 280: RPM UDP Parameters *(Continued)*

Parameter	Description
	searching for a Test Agent associated with the device. You can also hover over an interface name to view the details such as the Test Agent name, bind interface, bind family, and to check if the Test Agent is run on a device or not. You can also remove the interface filters you have selected and reset the filters to the default mode by using the Reset filters button.
Servers*	Specify the remote IP address or hostname of the server to which Test Agents send the requests. When you click the text box, the Select server page appears where you can enter details. To add more than one servers, click + Add Server and specify the following: • Servers—Specify the hostname or the IP address of the remote endpoint. Maximum length—255 characters. • Remote TCP port—Specify the port number of UDP. Default value—7. Range—7 through 65535.
Time Between Requests	Specify the time taken between successive UDP requests initiated by the routers. Unit—Seconds (s). Default value—10 s. Range—1 s through 255 s.

Table 280: RPM UDP Parameters (Continued)

Parameter	Description
Remote port	Configure the remote port number for the Test sessions. Default value—7. Range—7 through 65535.
Advanced	
Collection Interval	Specify the collection interval. Collection interval is the frequency at which the results are collected from the remote endpoint or the device. Collection interval must be larger than the specified Time between requests. Unit—Seconds (s). Default value—15 s. Range—5 s through 300 s.
Device Response Timeout	Specify the device response timeout value. Device response timeout determines the time to wait for a response from the router. If the specified time is exceeded, the connection to the device is considered failed. Device response timeout must be larger than the specified Collection Interval. Unit—Seconds (s). Default value—200 s. Range—30 s through 300 s.
Routing instance	Specify the name of routing instance on the router used by the probes. Maximum value—64.

Table 280: RPM UDP Parameters *(Continued)*

Parameter	Description
Data Size	Specify the size of the data portion of the request packets sent from the router. Default value—0. Range—0 through 65400.
Data fill	Specify the contents of the data portion of the request packets sent from the router. The value should be in hexadecimal format. The value should be in hexadecimal format.
TTL	Specify the number of times the packets hop before a device discards the packet. Time to live (TTL) value indicates lifespan of a request packet. TTL prevents infinite loop in a network when the packet fails to reach the destination. For example, if you have entered the TTL value as 64, every time the packet passes a device, the value is reduced by one until the packet reaches the remote server. If the entered value reaches zero before reaching the remote endpoint, the packet is discarded. Default value—64. Range—1 through 254.
DSCP Code Points	Specify the Differentiated Services Code Point (DSCP) value that is used in the IP packet headers. Maximum length—64 characters.
Hardware Timestamp	Enable hardware timestamping of RPM probe messages in the Packet Forwarding Engine host processor on the router. By default, the toggle button is disabled.

Table 280: RPM UDP Parameters *(Continued)*

Parameter	Description
Ping Timestamp	Enable ping timestamping to send UDP timestamp requests. By default, the toggle button is disabled.
One Way Hardware Timestamp	Enable the one-way hardware timestamping for one-way delay and jitter measurements. By default, the toggle button is disabled.
IPv6 Local Link	Specify the router link-local logical interface name for the egress interface to be used. This applies only for Test sessions with IPv6 address as the target address. Maximum length—64 characters.

Table 281: Y.1731 DM Parameters

Parameter	Description
General	

Table 281: Y.1731 DM Parameters (Continued)

Parameter	Description
Senders*	Select one or more interfaces of the network devices on which you want to run a Test. To select one or more network devices as client values: 1. Click the Clients text box. The Test Agent/Interfaces choices page appears. 2. Select one or more Test Agent interfaces. For information on filtering Test Agent interfaces, see "Search for Test Agent interfaces by using Filters" on page 1083. 3. Click OK. NOTE: On the Test Agent/Interfaces choices page, you can filter the Test Agent interfaces based on the following criteria: • Interface name or the Test Agent name, or both. Based on your requirement, you can select or deselect Test Agents or Interfaces. NOTE: The Test Agent filter is not functional as the RPM plugins directly interact with the devices. • IPv4 address (with or without IPv4 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv4 or No IPv4 to view or hide the interfaces that are associated with the IPv4 family and the interfaces that are not associated with the IPv4 family. • IPv6 address (with or without IPv6 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv6 or No IPv6 to view or hide the

Table 281: Y.1731 DM Parameters *(Continued)*

Parameter	Description
	interfaces that are associated with the IPv6 family and the interfaces that are not associated with the IPv6 family. - Deployment type of a Test Agent: - Application—Filter out the Test Agents deployed in a device as an application. - Device—Filter out the Test Agents associated with a device. Based on your requirement, you can select or deselect Application or Device to view or hide the Test Agents deployed in a device as an application or the Test Agents associated with a device. - Connection types between the interface and the router: - Link—Indicates that the status is Up and there is an active connection for the interface. - No Link—Indicates that the status is Down and there is no active connection for the interface. Based on your requirement, you can select or deselect Link or No Link. - Connection statuses of the Test Agent: - Online—The Test Agent is connected to Routing Director. - Offline—The Test Agent is not connected to Routing Director.

Table 281: Y.1731 DM Parameters *(Continued)*

Parameter	Description
	When the interface supports RPM plugins, the interface contains <i>(On device)</i> at the end of its name indicating the Device Online status. When the interface supports only the normal plugins, the interface will not display <i>(On device)</i> at the end of its name, and the status refers to the Online status. Based on your requirement, you can select or deselect Online or Offline. • Interface types of the Test Agent: • Normal—Network interfaces that are used for testing traffic. • Management—Network interfaces used by the Test Agents to connect to Routing Director. • Unknown—Network interfaces that the Test Agent is unable to identify or that have been removed. Based on your requirement, you can select or deselect Normal, Management, or Unknown. You can enable the Display Device Names toggle button to display the hostnames of the routers that are associated with the Test Agents. You can also enable the Hide unsupported toggle button to hide devices that are not supported by the plugin you have selected. To search for a specific interface or a Test Agent, enter one or more keywords in the search text box. You can search based on the name, description, IP and MAC address of the interfaces or the Test Agents. You can also search based on the device name, device model, and device MAC addresses if

Table 281: Y.1731 DM Parameters *(Continued)*

Parameter	Description
	searching for a Test Agent associated with the device. You can also hover over an interface name to view the details such as the Test Agent name, bind interface, bind family, and to check if the Test Agent is run on a device or not. You can also remove the interface filters you have selected and reset the filters to the default mode by using the Reset filters button. NOTE: The filters IPv4 and IPv6 are not applicable in the case of Y.1731 SLM.

Table 281: Y.1731 DM Parameters (Continued)

Parameter	Description
SLA Iterators	Specify one or more performance monitoring iterator profile names. When you click the text box, the Select iterator page appears where you can enter details. To add more than one iterator, click + Add iterator and specify the following: • Iterator profile name—Specify the iterator profile name used for SLA performance monitoring. • Maintenance domain name—Specify the maintenance domain name. A maintenance domain is a unique identifier for managing a network within a specific boundary. The domain name represents the space that is owned and operated by a single entity, which is defined by the set of internal ports and boundary ports. • Maintenance association name—Specify the maintenance association name. A maintenance association identifies services within a maintenance domain, and it logically groups the endpoints involved in monitoring the same service. The CMF protocol operates within this association. • Local maintenance end point—Specify the local maintenance endpoint (MEP). A local MEP is a device or a port within a maintenance domain. It is responsible for initiating a measurement to the remote MEP as well as detecting issues within a maintenance domain. • Remote maintenance end point—Specify the local maintenance endpoint (MEP). A remote MEP is a device or a port within a maintenance domain. It is responsible for initiating

Table 281: Y.1731 DM Parameters *(Continued)*

Parameter	Description
	a measurement to the local MEP as well as detecting issues within a maintenance domain. • Name—Specify a unique identifying name for the SLA iterator.
Collection interval (s)	Specify the collection interval. Collection interval is the frequency at which the results related to delay and jitter are collected from the router. Unit—Seconds (s). Default value—15 s. Range—5 s through 300 s.
Device response timeout	Specify the device response timeout value. Device response timeout determines the time to wait for a response from the router. If the specified time exceeds then the connection to the router is considered to be failed. Device response timeout must be larger than the specified Collection Interval. Unit—Seconds (s). Default value—200 s. Range—30 s through 300 s.

Table 282: Y.1731 SLM Parameters

Parameter	Description
General	

Table 282: Y.1731 SLM Parameters *(Continued)*

Parameter	Description
Senders*	Select one or more interfaces of the network devices on which you want to run a Test. To select one or more network devices as client values: 1. Click the Clients text box. The Test Agent/Interfaces choices page appears. 2. Select one or more Test Agent interfaces. For information on filtering Test Agent interfaces, see "Search for Test Agent interfaces by using Filters" on page 1089. 3. Click OK. NOTE: On the Test Agent/Interfaces choices page, you can filter the Test Agent interfaces based on the following criteria: • Interface name or the Test Agent name, or both. Based on your requirement, you can select or deselect Test Agents or Interfaces. NOTE: The Test Agent filter is not functional as the RPM plug-ins directly interact with the devices. • IPv4 address (with or without IPv4 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv4 or No IPv4 to view or hide the interfaces that are associated with the IPv4 family and the interfaces that are not associated with the IPv4 family. • IPv6 address (with or without IPv6 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv6 or No IPv6 to view or hide the

Table 282: Y.1731 SLM Parameters *(Continued)*

Parameter	Description
	interfaces that are associated with the IPv6 family and the interfaces that are not associated with the IPv6 family. - Deployment type of a Test Agent: - Application—Filter out the Test Agents deployed in a device as an application. - Device—Filter out the Test Agents associated with a device. Based on your requirement, you can select or deselect Application or Device to view or hide the Test Agents deployed in a device as an application or the Test Agents associated with a device. - Connection types between the interface and the router: - Link—Indicates that the status is Up and there is an active connection for the interface. - No Link—Indicates that the status is Down and there is no active connection for the interface. Based on your requirement, you can select or deselect Link or No Link. - Connection statuses of the Test Agent: - Online—The Test Agent is connected to Routing Director. - Offline—The Test Agent is not connected to Routing Director.

Table 282: Y.1731 SLM Parameters *(Continued)*

Parameter	Description
	When the interface supports RPM plugins, the interface contains <i>(On device)</i> at the end of its name indicating the Device Online status. When the interface supports only the normal plugins, the interface will not display <i>(On device)</i> at the end of its name, and the status refers to the Online status. Based on your requirement, you can select or deselect Online or Offline. • Interface types of the Test Agent: • Normal—Network interfaces that are used for testing traffic. • Management—Network interfaces used by the Test Agents to connect to Routing Director. • Unknown—Network interfaces that the Test Agent is unable to identify or that have been removed. Based on your requirement, you can select or deselect Normal, Management, or Unknown. You can enable the Display Device Names toggle button to display the hostnames of the routers that are associated with the Test Agents. You can also enable the Hide unsupported toggle button to hide devices that are not supported by the plugin you have selected. To search for a specific interface or a Test Agent, enter one or more keywords in the search text box. You can search based on the name, description, IP and MAC address of the interfaces or the Test Agents. You can also search based on the device name, device model, and device MAC addresses if

Table 282: Y.1731 SLM Parameters *(Continued)*

Parameter	Description
	searching for a Test Agent associated with the device. You can also hover over an interface name to view the details such as the Test Agent name, bind interface, bind family, and to check if the Test Agent is run on a device or not. You can also remove the interface filters you have selected and reset the filters to the default mode by using the Reset filters button. NOTE: The filters IPv4 and IPv6 are not applicable in the case of Y.1731 SLM.

Table 282: Y.1731 SLM Parameters *(Continued)*

Parameter	Description
SLA Iterators*	Specify one or more performance monitoring iterator profile names. When you click the text box, the Select iterator page appears where you can enter details. To add more than one iterator, click + Add iterator and specify the following: • Iterator profile name—Specify the iterator profile name used for SLA performance monitoring. • Maintenance domain name—Specify the maintenance domain name. A maintenance domain is a unique identifier for managing a network within a specific boundary. This domain name represents the space that is owned and operated by a single entity, which is defined by the set of internal ports and boundary ports. • Maintenance association name—Specify the maintenance association name. A maintenance association identifies services within a maintenance domain, and it logically groups the endpoints involved in monitoring the same service. The CMF protocol operates within this association. • Local maintenance end point—Specify the local maintenance endpoint (MEP). A local MEP is a device or a port within a maintenance domain. It is responsible for initiating a measurement to the remote MEP as well as detecting issues within a maintenance domain. • Remote maintenance end point—Specify the local maintenance endpoint (MEP). A local MEP is a device or a port within a maintenance domain. It is responsible for initiating

Table 282: Y.1731 SLM Parameters *(Continued)*

Parameter	Description
	a measurement to the remote MEP as well as detecting issues within a maintenance domain. • Name—Specify a unique identifying name for the SLA iterator.
Collection interval (s)	Specify the collection interval. Collection interval is the frequency at which the results related to delay and jitter are collected from the router. Unit—Seconds (s). Default value—15 s. Range—5 s through 300 s.
Device response timeout	Specify the device response timeout value. Device response timeout determines the time to wait for a response from the router. If the specified time exceeds then the connection to the router is considered to be failed. Device response timeout must be larger than the specified Collection Interval. Unit—Seconds (s). Default value—200 s. Range—30 s through 300 s.

Configure Evaluation Criteria

Table 283: Evaluation Criteria

Field	Description
Field	Select the type of metric that you want to evaluate from the drop-down list. The metrics listed in this drop-down depends on the Task you select. The metrics will be displayed as stream graphs in the <i>Stream-Name</i> Details page.
Comparator	Select the type of comparator that you want to use for the evaluation. You can choose among the following comparators— ==(equal to), != (not equal to), < (less than), <= (less than or equal to), > (greater than), >= (greater than or equal to). For example, if you have used > (greater than) comparator, and if you have specified 3000 ms in the Value text box, then an event is raised if the value exceeds 3000 ms.
Value	Specify the threshold value that you want to evaluate. On violation of any configured threshold value, an event is generated.

Table 283: Evaluation Criteria (*Continued*)

Field	Description
Severity	Select the severity level of the event that is raised when there is a threshold violation. The severity levels are: Debug—Indicates that the event needs troubleshooting. Info—Indicates that it is an informational event and does not require immediate attention. Warning—Indicates that the event needs to be reviewed but does not require immediate attention. Error—Indicates that the event needs to be reviewed and requires immediate attention. Critical—Indicates that the event is major and needs immediate attention. When an event is raised, the severity level you have selected will be displayed in the Events bar of the <i>Test-Name</i> page.
Name	Specify the name for the evaluation criteria.
Description	Specify the description to be displayed when there is a threshold violation. If you do not specify the description, the default description generated by Routing Director is displayed. When an event is raised, the description you have specified here will be displayed in the Events bar of the <i>Test-Name</i> page.
Raise Delay	Specify the duration that the Routing Director waits before triggering an event. If the number of seconds exceeds the configured Raise delay value, then an event is triggered. The triggered event is displayed in the Events bar of the <i>Test-Name</i> page.

Table 283: Evaluation Criteria *(Continued)*

Field	Description
Clear Delay	Specify the duration of time after which an event is cleared if the threshold is not violated. If the threshold has not exceeded the configured Clear delay value, then the event is cleared.

RELATED DOCUMENTATION

Active Assurance Workflow 891
About the Tests Page 1100

Create Input Variables

IN THIS SECTION

- Create Input Variables | 1097
- Reuse Input Variables | 1098
- Manage input variables | 1098

Create Input Variables

Input variables are reusable parameters that you define when you configure a Step for a Test or a Test template. Input variables ensures that the same values are consistently used across different Steps, thereby eliminating the need to reconfigure the variables each time and minimizing the risk of discrepancies.

To create an input variable for a parameter:

1. Configure the Test as described from Step 1 ["Step 1" on page 958](#) through Step 4 ["Step 4" on page 960](#) of ["Create a Test" on page 958](#).
2. Click the **Link** icon next to the parameter field for which you want to create an input variable.

A Create an Input Variable page appears.

3. Specify the details for the input variable as described in [Table 284 on page 1099](#).



NOTE: Fields marked with an asterisk (*) are mandatory.

4. Click **Save** to save the input variable.

The input variable you created is displayed as a *Label-Name* tag in the parameter field you are configuring.

Reuse Input Variables

To reuse input variables, you must first create input variables as described in ["Create Input Variables" on page 1097](#). You can reuse these existing variables when configuring any new Tasks and Steps.

1. At the time of Test creation, select the **Link** icon next to the parameter field for which you want to use an existing input variable.

A Select Variable page appears.

2. Click the drop-down list icon in the **Existing variables** field.
3. Select the desired variable from the list.
4. Click **OK**.

The input variable you selected is displayed as a *Label-Name* tag in the parameter field you selected.

Manage input variables

Once you create input variables for all the Steps, you can manage the input variables before executing the Test.

To manage input variables for a Test, click the **Global** icon on the right side of the Measurement Designers page and a Input Variables page appears with two tabs:

- **Manage tab**—In this tab, you can manage the list of all the input variables you have defined for a Test. It gives you an overall view of variables ensuring consistency across the Test configurations. In the Manage tab, there are two sections:
 - **Source section**—In this section, you can:
 - **View input variables**—View the list of input variables with their label, type, and description.

- **Create new input variable**—Create a new input variable by clicking the Add (+) icon and following the steps described in ["Create Input Variables" on page 1097](#).
- **Delete an input variable**—Select and delete an input variable from the list. To delete an input variable, select the input variable that you want to delete and click the **Delete** (trash can) icon. A message is displayed asking for your confirmation. Click **Yes** and the input variable is deleted.
- **Target section**—In this section, you can view the Steps, Tasks, and Parameters where you have applied the input variables listed in the Source section. You can also assign available input variables to different Steps, Tasks, and Parameters.
- **Preview tab**—In this tab, you can edit the Test details and input variable details before executing a Test. It helps to dynamically update the parameters giving flexibility in Test execution. In the Preview tab, you can:
 - **Add or edit Test details**—You can add or edit the name and tags of the Test. You can update or add these details as needed.
 - **Add or edit input variables**—You can add values or edit the existing input variables configured while creating a Test.

Table 284: Fields in the Create an Input Variable page

Fields	Descriptions
Type*	By default, the Type field is pre-populated. You cannot edit this field.
Label*	By default, the Label field is pre-populated. However, you can edit the existing label name or specify a new label name for the parameter.
Variable Name*	By default, the Variable Name field is pre-populated. However, you can edit the existing variable name or specify a new variable name for the parameter.
Group	Assign the variable to a group for better organization.
Unit	Specify the unit for the variable.

Table 284: Fields in the Create an Input Variable page *(Continued)*

Fields	Descriptions
Description	Specify a description for the variable.
Allow Multiple	Enable the Allow Multiple toggle button to allow a Test to use multiple Test Agent interfaces. The Allow Multiple field is displayed only if the specified Clients parameters involve interfaces.
Maximum	Set the maximum allowable value for the selected variable.
Minimum	Set the minimum allowable value for the selected variable.
Plugin	By default, the Plugin field displays the plug-in on which you have configured the variable.

RELATED DOCUMENTATION

| [Active Assurance Workflow](#) | 891

About the Tests Page

IN THIS SECTION

● [Tasks You Can Perform](#) | 1101

To access this page, click **Observability > Active Assurance > Tests**.

Tests are entities that you configure to measure metrics in a time-bound manner to produce a static measurement for the given duration. You (superusers and network administrators) can use the Tests

page to view a list of Tests that the Test Agents run in your network and the details of these Tests. For more information, see ["Tests and Monitors Overview" on page 944](#).

The widgets display the following information:

- **Total**—The total number of Tests that you have run in your network.
- **Passed**—The number of Tests completed successfully. Routing Director considers a Test as passed when the KPIs have not exceeded the evaluation criteria you have specified. Routing Director also displays the pass percentage of Tests.
- **Failed**—The number of Tests that failed. Routing Director considers a Test as failed when any one of the KPIs have exceeded the evaluation criteria you have specified. Routing Director also displays the fail percentage of Tests.
- **Error**—The number of Tests that have encountered errors. Routing Director also displays the error percentage of Tests. An error occurred can be of any reasons from a Test Agent being offline to a metric configuration failure or a timeout.



NOTE: By default, these widgets display the data for the past 1 week. However, you can customize the data visualization for a specific time range. For example, if you click **2h**, the widgets displays the Tests that were run for the past 2 hours.

Tasks You Can Perform

You can perform the following tasks on the Tests page:

- **View details of a Test**—You can view the list of all the Tests that you have run and the details of Tests. To view the details of a Test, select a Test and click **More > Details**. You can also hover over the *Test-Name* and click the **Detailed View** icon. The Test execution details pane appears on the right side of the page displaying the Test details.

On the Test execution details pane, you can:

- **Copy API Request URL**—You can copy the API Request URL so that you can reuse the copied API request URL to fetch the Test details when you rerun the Test.
- **Download Results**— You can download the Test result summary as a JSON file to your local system. When you click the **Download Results** button, a JSON file is generated that you can download to your local system.

The JSON file includes information displayed in the Tests execution details pane that can be used to analyze the results locally.

- You can view additional information such as, Test ID, Test name, Test status (passed or failed), and so on.

Click the Close (x) icon to close the pane.

- View details of a selected Test. See ["About the Test-Name Page" on page 1104](#).
- Create a Test. See ["Create a Test" on page 958](#).
- You can also perform the following tasks on this page:
 - Sort, resize, or re-arrange columns in a table (grid).
 - Show or hide columns in the table or reset page preferences, using the vertical ellipsis menu.
 - Filter the data displayed in the table—Click the filter icon (funnel) and select whether you want to show or hide advanced filters. You can then add or remove filter criteria, save criteria as a filter, apply or clear filters, and so on. The filtered results are displayed on the same page.

For more information, see ["GUI Overview" on page 6](#).

Table 285: Fields on the Tests Page

Fields	Description
Name	The name of the Test that you specified when you created the Test.

Table 285: Fields on the Tests Page (Continued)

Fields	Description
Status	The status of the Test: • Scheduled—The Test is scheduled to run. • Waiting—The Test Agent is getting ready to run the Test. • Running—The Test is running. • Stopping—The Test is being stopped and preparing the results. • Error—The Test has encountered an error while being executed by the Test Agent. The errors can be anything from a Test Agent being offline to a metric configuration failure or a timeout. • Passed—The Test has passed successfully. • Failed—The Test has failed because values for some metrics exceed the threshold you configured. • Canceled—The Test has been canceled by you. Routing Director retains the data related to the canceled Test up to the point of cancellation.
Status Message	The status message generated by Routing Director after you run a Test.
Execution Start Time	The date and time when the Test was executed. The timestamp is displayed in the following format: Month DD, YYYY, HH:MM:SS AM/PM. For example, Feb 5, 2024, 4:29:52 PM.
Execution End Time	The date and time when the Test ended. The timestamp is displayed in the following format: Month DD, YYYY, HH:MM:SS AM/PM. For example, Feb 5, 2024, 4:29:52 PM.
Executed By	The username (e-mail address) of the user who has run the Test.

Table 285: Fields on the Tests Page (*Continued*)

Fields	Description
Test Tags	The tags you have configured for the Test in the <i>key:value</i> format. By default, only two tags are displayed in the column. To view more than two tags, click + Number more. A tag is a key-value pair in which the key signifies a category for which you configure a value. The value is an identifier for the category. Examples of key-value pairs are device and device name (edgedevice:acx7000), site and site name (site:bangalore). You can configure tags for a Test at the time of creating a Test.

RELATED DOCUMENTATION

[About the Test Agents Page](#) | 900

About the *Test-Name* Page

IN THIS SECTION

● [Tasks You Can Perform](#) | 1105

To access the *Test-Name* page:

1. Select **Observability > Active Assurance > Tests**.

The Tests page appears.

2. Click a *Test-Name*.

The *Test-Name* page appears displaying the Test results.

You (superusers and network administrators) can use this page to view the details of the Test you run. You can view general details of a Test such as the Test name, Test description, date and time of the

completed Test, and the overall result (Passed, Failed, or Canceled), and the list of events generated. You can also view the details of Tasks generated for each Step, and Streams generated for each Task that you have configured for a Test. In addition, you can also cancel, rerun, clone, and delete a Test.

Tasks You Can Perform

You can perform the following Tasks on the *Test-Name* page:

- View Test details—You can view the following details related to the Test:
 - Test name—The name of the Test that you have specified at the time of creating the Test.
 - Test description—The description of the Test that you have entered at the time of creating the Test.
 - Full Test Execution duration—The total duration of the Test.
 - Test creation time—The date and time when the Test was created.
 - Test status—The overall status of the Test that you run.
- View the details of each Step within the Test—A Test can include multiple Steps. Each Step within a Test is displayed as a tab. Click a Step (tab) to view the Step results, the details of the Tasks within the Step, events generated for the Step, and the Streams for each Task within the Step. You can also view the time duration up to which Routing Director waits for Test Agent to execute the Test and the total run time of Test.

If a single Step fails, Routing Director evaluates the overall Test result status as Failed.

- View the list of events—You can view the events generated by the Test and the date and time at which the event has occurred. The Event bar displays events based on the severity levels. The high-priority event, **Critical**, is displayed at the top of the events list, while low-priority event, **Information**, is displayed at the end of the list.

To view the list of all the events generated in the order of occurrence, click the **More** option. When you click **More**, the Events page appear. For more information on the events generated by the Test, see [Table 286 on page 1107](#).

- View the events on an event bar—You can view all the events generated on an event bar. The event bar is a color-driven bar graph that indicates events and the time at which the event has occurred. The colors represent the severity of events. That is, critical events are represented in red, errors are represented in orange, warnings are represented in yellow, information is represented in blue.

You can also hover over an event generated on the event bar to view the start time, end time, total number of events generated, and the number of events generated in each type.

- View the status of each Task within a Step—A Step is a collection of Tasks. Each Task within a Step is displayed as different event bars. Click > to view the event bar for each Task. When you hover over the event bar, you can view the events generated for the Task.
- View the list of Streams—You can view the list of Streams generated for the Task you selected and their details as displayed in [Table 287 on page 1108](#). To view the details of a specific Stream, click a *Stream-Name*. The *Stream-Name* Details page displays the result for metrics as stream graphs. For more information, see ["View Stream Details" on page 1258](#).
- Group all the Tasks—By default, all Streams are displayed in a table format. When you enable the **Group by Task** toggle button, the Streams are grouped based on the Task that you have selected. For example, all the Streams related to HTTP will be grouped under HTTP Task. By default, this toggle button is disabled.

When you enable this toggle button, you can expand and hide the Streams within the Tasks by using **Expand** and **Minimize** icons. When you click **Expand** icon, the Stream details for each Task are displayed, and when you click **Minimize**, the Stream details for each Task are hidden.

- Cancel the Test—You can cancel a Test by using the Cancel option. To cancel a Test, click **More** and select **Cancel**. A message is displayed asking for your confirmation. Click **Yes**.

A message confirms that Routing Director has successfully canceled the Test, and the Test status is changed to **Canceled**. When you cancel a running Test, Routing Director retains the data related to the canceled Test up to the point of cancellation.

- Rerun a Test—You can rerun a Test by using the Rerun option. To rerun a Test, click **More** and select **Rerun**.

A message confirms that Routing Director successfully rerun the Test and the overall status of the Test changes to Running.

- Clone a Test—You can clone a Test by using the Clone option. To clone a Test, click **More** and select **Clone**. When you clone a Test, an editable copy of the Test you configured is displayed.

You are redirected to the Measurement Designer (**Observability > Active Assurance > Measurement Designers**) page and a clone of the Test you configured is displayed.

- Delete the Test—You can delete a Test by using the Delete option. To delete the Test, click **More** and select **Delete**. A message is displayed asking for your confirmation. Click **Yes**.

A message confirms that Routing Director has successfully deleted the Test and you are redirected to the Tests (**Observability > Active Assurance > Tests**) page. Routing Director also deletes the Test from the Tests page and the data on the widgets is updated.

Table 286: Fields in the Events page

Field	Description
Severity	The type of severity level of the event that is raised when there is a criteria violation. The following are the severity levels: Critical—Indicates that the event is critical and needs immediate attention. Warning—Indicates that the event needs to be fixed but does not require immediate attention. Info—Indicates the progress of the Task and provides information on the event. The informational event does not require attention. Error—Indicates that the event needs to be fixed and requires immediate attention and troubleshooting.
Description	The description that you specified when you configured the evaluation criteria for a Test.
Raise Time	The date and time when the event was generated. The timestamp is displayed in the following format: Month DD, YYYY, HH:MM:SS AM/PM. The date and time is displayed according to the Raise delay you specified when you configured the evaluation criteria for a Test. For example, Mar 5, 2024, 4:29:52 PM.
Clear Time	The date and time when the event was cleared. The timestamp is displayed in the following format: Month DD, YYYY, HH:MM:SS AM/PM. The date and time is displayed according to the clear delay you specified when you configured the evaluation criteria for a Test. For example, Mar 5, 2024, 4:29:52 PM.

Table 286: Fields in the Events page *(Continued)*

Field	Description
Subject	The details of the event generated is displayed in JSON format. It displays various IDs related to the event like Test Agent ID, Test ID, Task ID, Stream ID, and so on. To view the details, click show hyperlink.
Data	The details of the evaluation criteria associated with the event generated. To view data, click show hyperlink.

Table 287: Fields in the Test Result Table

Field	Description
Stream Name	The name of the Stream that is auto-generated by Routing Director when you run a Test.
Config	The configuration details of the selected Stream. Click show to view the details of the results metrics measured by the Stream. You can also copy the configuration.
Task Name	The name of the Task that generated the Stream. Tasks such as a ping are bi-directional. A bi-directional Task produces more than one Stream in each direction.

Table 287: Fields in the Test Result Table *(Continued)*

Field	Description
Severity	The severity level of events raised by Streams. The following are the severity levels: Debug—Indicates that the event needs troubleshooting. Critical—Indicates that the event is critical and needs immediate attention. Warning—Indicates that the event needs to be fixed but does not require immediate attention. Info—Indicates that an informational event is raised that provides details on the progress of the Task and does not require attention. Error—Indicates that the event needs to be fixed and requires immediate attention and troubleshooting.
Event Overview	The event overview of a Stream displays an event bar. You can view the start time, end time, total number of events generated, and the number of events generated in each type. The events can have Critical (Red), Error (Orange), Warning (Yellow), or Info (Blue) status. If you click a <i>Stream-Name</i>, you can access a detailed result of all metrics measured by the Stream. See "View Stream Details" on page 1258 for more information.

RELATED DOCUMENTATION

Tests and Monitors Overview | 944

Create a Monitor

IN THIS SECTION

- [Configure Monitor Settings | 1110](#)
- [Configure Task Parameters | 1113](#)
- [Native Plug-ins | 1115](#)
- [RPM Plug-ins | 1198](#)
- [Configure Evaluation Criteria | 1245](#)

Configure Monitor Settings

A Monitor is a set of verifications that are performed by one or more Test Agents for an infinite amount of time. A Monitor contains a single step with one or more parallel Tasks, which continuously monitors the KPIs that you have defined in the Monitor. A Task contains the configuration to measure specific metrics. Monitors measure the metrics indefinitely until you decide to stop the Monitor. For more information on Monitors, see ["Tests and Monitors Overview" on page 944](#).

To create a Monitor:

1. Do one of the following to access the Measurement Designer page:
 - a. Access Measurement Designer page through the Monitors page.
 - i. Navigate to the Monitors (**Observability > Active Assurance > Monitors**) page.
The Monitors page appears.
 - ii. Click the **Add (+)** icon.
The Create new page appears.
The Create new page displays **All**, **Tests**, and **Monitors** mode. By default, the Monitors mode is enabled.
 - iii. Select the **+ Create blank Monitor**.
The Measurement Designer page appears.
 - b. Access Measurement Designer page directly.

- i. Navigate to the Measurement Designer (**Observability > Active Assurance > Measurement Designer**) page.

The Create new page appears.

The Create new page displays **All**, **Tests**, and **Monitors** mode, and the All mode is enabled by default.

- ii. Select the **Monitors** mode.
- iii. Select the **+ Create blank Monitor**.

The Measurement Designer page appears.



NOTE: You can also access Monitor creation mode from the Test creation mode by using the Monitor button. If you want to access Monitor creation mode, ensure that you have configured only one Step. If you have configured more than one Step, the Steps other than the first one will not be moved to the Monitor creation mode.

2. Specify a name for the Monitor.

Click the **Monitor Name** or click the **Edit** (pencil) icon to enter a name for the Monitor in the **Monitor Name** text box

Monitor Name is the placeholder name for the Monitor.



NOTE: You cannot run a Monitor if you have not specified the **Monitor Name**.

3. Select a Task from the Tasks page displayed by default, or click **+ Add Task** to add Tasks for a Monitor. The Tasks page displays the list of protocols that can be configured for a Monitor.

In the Tasks page, do one of the following:

- Select a Task from the Tasks page. For example, click DNS to add DNS as one of the Tasks in the Step.
- Click the **Add (+)** icon next to the Task.
- Drag and drop the Task inside the Step that you are configuring.

Select the Close (X) icon to close the Tasks page.

For a Task, you can do the following:

- (Optional) Edit the name of a Task—To edit a Task, do one of the following:

- Click the **Edit** (pencil) icon to specify a name for the Task in the *Task-Name* text box.
- Click the *Task-Name* and specify a name for the Task in the *Task-Name* text box.

By default, the plugin name is displayed. If you do not edit, the default name will be used.

- (Optional) Delete the Task—To delete a Task, do one of the following:
 - Click the horizontal ellipsis in the Task box and click **Remove**.
 - Drag the Task anywhere on the empty portion of the screen the screen.

A confirmation message appears asking to confirm if you want to remove the Task. Click **Yes**.

- Configure the parameters for a Task—Click the **Settings** (gear) icon on the Task box to configure a Task. The **Monitor** page appears and displays the Task you added. This page includes the following tabs:
 - Parameters tab—Configure parameters for the Tasks that you have added. For more information on parameters that you can configure for a Task, see [Table 288 on page 1113](#).
 - Evaluation criteria tab—The Evaluation Criteria for each Task is added by default. You can configure customized evaluation criteria for metrics by using threshold expressions. On violation of any configured expression, Routing Director generates events. For more information on metrics that you can configure, see [Table 307 on page 1245](#).



NOTE: Fields marked with an asterisk (*) are mandatory.

4. A Monitor can contain only one Step, and this Step can include one or more Tasks. Based on your requirements, repeat "[Step 3" on page 1111](#) to add one or more Tasks.

5. Click **Monitor Settings** (gear) icon on the right side of the page to,

- (Optional) Specify a short description for the Monitor.
- (Optional) Add **Tags**—Specify a key-value pair in the key:value format.

You can configure Tags for the Monitor in the key:value format to provide additional information about the Monitor you are configuring. A tag is a key-value pair in which the key signifies a category for which you configure a value. The value is an identifier for the category. Examples of key-value pairs are device and device name (edgedevice:acx7000), site and site name (site:bangalore).

6. Click **Run**.

A message confirming that the Monitor is created successfully appears and you are redirected to the *Monitor-Name* (**Observability > Active Assurance > Monitor > *Monitor-Name***) page. On the *Monitor-*

Name page, Monitor status is displayed as **Running** indicating that the Monitor is in progress and you can view the details of the Monitor. See ["About the Monitor-Name page" on page 1104](#) for more information about the Monitor details.



NOTE: The **Run** button is disabled until you add a Task under a Step.

Configure Task Parameters

Table 288: Tasks and its parameters

Tasks	Description
DNS	For information on the parameters that you can configure for DNS, see Table 289 on page 1115 .
HTTP	For information on the parameters that you can configure for HTTP, see Table 290 on page 1122 .
Ping	For information on the parameters that you can configure for Ping, see Table 291 on page 1128 .
RPM CISCO TWAMP Reflector	For information on the parameters that you can configure for RPM Cisco TWAMP Reflector, see Table 292 on page 1135. NOTE: We recommend that you configure RPM Cisco TWAMP Reflector using API as configuring RPM Cisco TWAMP Reflector using the Routing Director GUI is a Beta feature.
IPTV MPEG	For information on the parameters that you can configure for IPTV MPEG, see Table 293 on page 1141 .
Netflix Speedtest	For information on the parameters that you can configure for Netflix Speedtest, see Table 294 on page 1147 .

Table 288: Tasks and its parameters *(Continued)*

Tasks	Description
OTT - HLS	For information on the parameters that you can configure for OTT-HLS, see Table 295 on page 1153 .
TCP	For information on the parameters that you can configure for TCP, see Table 296 on page 1159 .
UDP	For information on the parameters that you can configure for UDP, see Table 297 on page 1169 .
TWAMP/TWAMP Light	For information on the parameters that you can configure for TWAMP/TWAMP Light, see Table 298 on page 1181 .
TWAMP Reflector	For information on the parameters that you can configure for TWAMP Reflector, see Table 299 on page 1194 .
RPM HTTP	For information on the parameters that you can configure for RPM HTTP, see Table 300 on page 1198 .
RPM PING	For information on the parameters that you can configure for RPM PING, see Table 301 on page 1204 .
RPM TCP	For information on the parameters that you can configure for RPM TCP, see Table 302 on page 1211 .
RPM TWAMP	For information on the parameters that you can configure for RPM TWAMP, see Table 303 on page 1217 .
RPM UDP	For information on the parameters that you can configure for RPM UDP, see Table 304 on page 1225 .

Table 288: Tasks and its parameters *(Continued)*

Tasks	Description
Y.1731 DM	For information on the parameters that you can configure for Y.1731 DM, see Table 305 on page 1232 .
Y.1731 SLM	For information on the parameters that you can configure for Y.1731 SLM, see Table 306 on page 1238 .

Native Plug-ins

Table 289: DNS Parameters

Parameter	Description
General	

Table 289: DNS Parameters (Continued)

Parameter	Description
Clients*	Select one or more Test Agent interfaces on which you want to run the Test. To select one or more Test Agents: 1. Click the Clients text box. The Select interfaces page appears. 2. Select one or more Test Agents. For information on filtering Test Agents, see "Search for Test Agent interfaces by using Filters" on page 1116. 3. Click OK. NOTE: On the Test Agent/Interfaces page, you can filter the Test Agent interfaces based on the following criteria: • Interface name or the Test Agent name, or both. Based on your requirement, you can select or deselect Test Agents or Interfaces. • IPv4 address (with or without IPv4 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv4 or No IPv4 to view or hide the interfaces that are associated with the IPv4 family and the interfaces that are not associated with the IPv4 family. • IPv6 address (with or without IPv6 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv6 or No IPv6 to view or hide the interfaces that are associated with the IPv6 family and the interfaces that are not associated with the IPv6 family. • Deployment type of a Test Agent:

Table 289: DNS Parameters *(Continued)*

Parameter	Description
	• Application—Filter out the Test Agents deployed in a device as an application. • Device—Filter out the Test Agents associated with a device. Based on your requirement, you can select or deselect Application or Device to view or hide the Test Agents deployed in a device as an application or the Test Agents associated with a device. • Connection types between the interface and the router: • Link—Indicates that the status is Up and there is an active connection for the interface. • No Link—Indicates that the status is Down and there is no active connection for the interface. Based on your requirement, you can select or deselect Link or No Link. • Connection statuses of the Test Agent: • Online—The Test Agent is connected to Routing Director. • Offline—The Test Agent is not connected to Routing Director. When the interface supports RPM plugins, the interface contains <i>(On device)</i> at the end of its name indicating the Device Online status. When the interface supports only the normal plugins, the interface will not display <i>(On</i>

Table 289: DNS Parameters (Continued)

Parameter	Description
	<i>device</i>) at the end of its name, and the status refers to the Online status. Based on your requirement, you can select or deselect Online or Offline. - Interface types of the Test Agent: - Normal—Network interfaces that are used for testing traffic. - Management—Network interfaces used by the Test Agents to connect to Routing Director. - Unknown—Network interfaces that the Test Agent is unable to identify or that have been removed. Based on your requirement, you can select or deselect Normal, Management, or Unknown. You can enable the Display Device Names toggle button to display the hostnames of the routers that are associated with the Test Agents. You can also enable the Hide unsupported toggle button to hide devices that are not supported by the plugin you have selected. To search for a specific interface or a Test Agent, enter one or more keywords in the search text box. You can search based on the name, description, IP and MAC address of the interfaces or the Test Agents. You can also search based on the device name, device model, and device MAC addresses if searching for a Test Agent associated with the device. You can also hover over an interface name to view the details such as the Test Agent name, bind interface, bind family, and to check if the Test Agent is run on a device or not.

Table 289: DNS Parameters *(Continued)*

Parameter	Description
	You can also remove the interface filters you have selected and reset the filters to the default mode by using the Reset filters button.
Lookup names*	Specify the domain name to lookup. For example, example.net Lookup is the process of querying the server to translate a domain name into IP address. When you click the Lookup names text box, a Select Lookup name page appears. You can add a name in the Lookup text box. To add more than one Lookup names, click + Add Lookup.
Time between requests	Specify the time that Test Agent should wait between successive DNS queries. Unit—seconds (s) Default value—10.00 s Range—0.01 s through 3600 s
DNS server	Specify the IP address of the DNS server to query. If left empty, Test Agents use the default DNS server. Maximum Length—200 characters

Table 289: DNS Parameters *(Continued)*

Parameter	Description
DNS record type	Select the DNS record type. A DNS record is a set of unstructured data stored in a DNS database. The database consists of information on a domain and its services. DNS has different resource records. Each record type has different functions in the resolution process. • A (Address)—Associates a domain name with an IPv4 address. • AAAA (IPv6 Address)—Associates a domain name with an IPv6 address. • SOA Record (Start of Authority)—Provides information about a DNS zone. • MX (Mail Exchange)—Associates a domain name to the email server responsible for receiving an email. • NS (Name Server)—Provides domain information of an address. It specifies the servers responsible for hosting DNS servers for a particular domain. • TXT (Text)—Provides storage for text data. It is also used for domain verification. • PTR (Pointer)—Performs reverse DNS operation. It associates an IP address with a domain name. • CNAME (Canonical Name)—Maps domain names to another domain and not an IP address. By default, the record type is A.
Thresholds for errored seconds (ES)	

Table 289: DNS Parameters (*Continued*)

Parameter	Description
Timeout	Specify the timeout value. If no response is received for the DNS request within the defined timeout period, an errored-second is raised unless the request lifetime expires. Unit—Milliseconds (ms) Default value—1000 ms Range—1ms through 30000 ms
Advanced	
Request lifetime	Specify the timeout value. If no response is received for the DNS request within the defined timeout period, an errored-second is raised unless the request lifetime expires. Request lifetime value is the duration for which a DNS request is alive. It determines how long a request persists without getting terminated. Unit—Milliseconds (ms) Default value—5000 ms Range—1 ms through 30000 ms
Recursive requests	Enable or disable the Recursive Requests toggle button. Recursive Request is the process where a DNS server queries other DNS servers until it resolves the lookup address. If you enable the Recursive Requests toggle button, the Recursion Desired flag is set on the DNS queries. By default, the toggle button is enabled.

Table 289: DNS Parameters *(Continued)*

Parameter	Description
Response code	Select the DNS response code. You can select one of the following response codes: • NOERROR—Indicates that the query was successful. • SERVFAIL—Indicates that the server has failed to complete the request. • NXDOMAIN—Indicates that the domain name does not exist. • REFUSED—Indicates that the server refused to perform the operation. • NOTAUTH—Indicates that the server is not authoritative for the zone. The DNS response code can be any value from 0 to 9. The value indicates the outcome of a DNS query. For more information on DNS response codes and strings with descriptions, see IANA link.
Expected response	Specify an expected response from the DNS server. If the actual response does not match the expected response, an errored-second is raised. In case, the response consists of more than one codes, at least one code must match the expected response. For more information on DNS response codes and strings with descriptions, see IANA link.

Table 290: HTTP Parameters

Parameter	Description
General	

Table 290: HTTP Parameters *(Continued)*

Parameter	Description
Clients*	Select one or more Test Agent interfaces on which you want to run the Test. To select one or more Test Agent interfaces: 1. Click the Clients text box. The Test Agent/Interfaces page appears. 2. Select one or more Test Agent interfaces. For information on filtering Test Agent interfaces, see "Search for Test Agent interfaces by using Filters" on page 1123. 3. Click OK. NOTE: On the Test Agent/Interfaces page, you can filter the Test Agent interfaces based on the following criteria: • Interface name or the Test Agent name, or both. Based on your requirement, you can select or deselect Test Agents or Interfaces. • IPv4 address (with or without IPv4 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv4 or No IPv4 to view or hide the interfaces that are associated with the IPv4 family and the interfaces that are not associated with the IPv4 family. • IPv6 address (with or without IPv6 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv6 or No IPv6 to view or hide the interfaces that are associated with the IPv6 family and the interfaces that are not associated with the IPv6 family.

Table 290: HTTP Parameters *(Continued)*

Parameter	Description
	• Deployment type of a Test Agent: • Application—Filter out the Test Agents deployed in a device as an application. • Device—Filter out the Test Agents associated with a device. Based on your requirement, you can select or deselect Application or Device to view or hide the Test Agents deployed in a device as an application or the Test Agents associated with a device. • Connection types between the interface and the router: • Link—Indicates that the status is Up and there is an active connection for the interface. • No Link—Indicates that the status is Down and there is no active connection for the interface. Based on your requirement, you can select or deselect Link or No Link. • Connection statuses of the Test Agent: • Online—The Test Agent is connected to Routing Director. • Offline—The Test Agent is not connected to Routing Director. When the interface supports RPM plugins, the interface contains <i>(On device)</i> at the end of its name indicating the Device Online status.

Table 290: HTTP Parameters *(Continued)*

Parameter	Description
	When the interface supports only the normal plugins, the interface will not display <i>(On device)</i> at the end of its name, and the status refers to the Online status. Based on your requirement, you can select or deselect Online or Offline. - Interface types of the Test Agent: - Normal—Network interfaces that are used for testing traffic. - Management—Network interfaces used by the Test Agents to connect to Routing Director. - Unknown—Network interfaces that the Test Agent is unable to identify or that have been removed. Based on your requirement, you can select or deselect Normal, Management, or Unknown. You can enable the Display Device Names toggle button to display the hostnames of the routers that are associated with the Test Agents. You can also enable the Hide unsupported toggle button to hide devices that are not supported by the plugin you have selected. To search for a specific interface or a Test Agent, enter one or more keywords in the search text box. You can search based on the name, description, IP and MAC address of the interfaces or the Test Agents. You can also search based on the device name, device model, and device MAC addresses if searching for a Test Agent associated with the device. You can also hover over an interface name to view the details such as the Test Agent name, bind

Table 290: HTTP Parameters *(Continued)*

Parameter	Description
	interface, bind family, and to check if the Test Agent is run on a device or not. You can also remove the interface filters you have selected and reset the filters to the default mode by using the Reset filters button.
URLs*	Specify the URL for the HTTP request. Maximum Length—200 characters
Time between requests	Specify the time taken between successive HTTP requests initiated by Test Agent. Unit—Seconds (s) Default value—10.00 s Range—0.01 s through 3600 s
Thresholds for errored seconds (ES)	
Timeout	Specify the timeout value. If no response is received to the HTTP request within the defined timeout period, an errored second for timeout is raised. Unit—Milliseconds (ms) Default value—1000 ms Range—1 ms through 30000 ms
Response content	Specify the regular expression (case-insensitive) against which the HTTP response content is validated. Maximum Length—50 characters
Advanced	

Table 290: HTTP Parameters *(Continued)*

Parameter	Description
Request lifetime	Specify the maximum time to wait for a response before the HTTP request is canceled. If this time expires without a response, an errored second for lifetime expiry is raised. Unit—Milliseconds (ms) Default value—5000 ms Range—1 ms through 30000 ms
HTTP response code	Specify the HTTP response code for Test Agent. If the HTTP response code does not match the specified response code, an errored-second is raised. For more information on HTTP response codes, see RFC 9110.
Proxy server	Specify the IP address of the HTTP proxy server optionally.
Proxy port	Specify the HTTP proxy port when using HTTP proxy server. Default value—8080 Range—1 through 65535

Table 290: HTTP Parameters *(Continued)*

Parameter	Description
Proxy authentication	Select the authentication method that the Test Agent uses when connecting to a proxy server. Select one of the following authentication method: • None—Indicates that you can access the remote endpoint without any authentication. • Basic—Indicates that you can access the remote endpoint by using a username and password. • Digest—Indicates that you can access the remote endpoint by using a username and password, but it will send a hashed version of the password, making the password resistant to attacks. • Ntlm—Indicates that you can access a remote endpoint by using Single Sign-on. That is, without a password. Proxy authenticates the incoming request from a Test Agent. This ensures that only the authorized users have access to the internet.
Proxy username	Specify the username for authorized access to a proxy.
Proxy password	Specify the password for authorized access to a proxy server.

Table 291: Ping Parameters

Parameter	Description
General	

Table 291: Ping Parameters (Continued)

Parameter	Description
Clients*	Select one or more Test Agent interfaces on which you want to run the Test. To select one or more Test Agent interfaces: 1. Click the Clients text box. The Test Agent/interfaces choices page appears. 2. Select one or more Test Agent interfaces. For information on filtering Test Agent interfaces, see "Search for Test Agent interfaces by using Filters" on page 1129. 3. Click OK. NOTE: On the Test Agent/interfaces choices page, you can filter the Test Agent interfaces based on the following criteria: • Interface name or the Test Agent name, or both. Based on your requirement, you can select or deselect Test Agents or Interfaces. • IPv4 address (with or without IPv4 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv4 or No IPv4 to view or hide the interfaces that are associated with the IPv4 family and the interfaces that are not associated with the IPv4 family. • IPv6 address (with or without IPv6 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv6 or No IPv6 to view or hide the interfaces that are associated with the IPv6 family and the interfaces that are not associated with the IPv6 family.

Table 291: Ping Parameters (Continued)

Parameter	Description
	• Deployment type of a Test Agent: • Application—Filter out the Test Agents deployed in a device as an application. • Device—Filter out the Test Agents associated with a device. Based on your requirement, you can select or deselect Application or Device to view or hide the Test Agents deployed in a device as an application or the Test Agents associated with a device. • Connection types between the interface and the router: • Link—Indicates that the status is Up and there is an active connection for the interface. • No Link—Indicates that the status is Down and there is no active connection for the interface. Based on your requirement, you can select or deselect Link or No Link. • Connection statuses of the Test Agent: • Online—The Test Agent is connected to Routing Director. • Offline—The Test Agent is not connected to Routing Director. When the interface supports RPM plugins, the interface contains (<i>On device</i>) at the end of its name indicating the Device Online status.

Table 291: Ping Parameters (Continued)

Parameter	Description
	When the interface supports only the normal plugins, the interface will not display (<i>On device</i>) at the end of its name, and the status refers to the Online status. Based on your requirement, you can select or deselect Online or Offline. - Interface types of the Test Agent: - Normal—Network interfaces that are used for testing traffic. - Management—Network interfaces used by the Test Agents to connect to Routing Director. - Unknown—Network interfaces that the Test Agent is unable to identify or that have been removed. Based on your requirement, you can select or deselect Normal, Management, or Unknown. You can enable the Display Device Names toggle button to display the hostnames of the routers that are associated with the Test Agents. You can also enable the Hide unsupported toggle button to hide devices that are not supported by the plugin you have selected. To search for a specific interface or a Test Agent, enter one or more keywords in the search text box. You can search based on the name, description, IP and MAC address of the interfaces or the Test Agents. You can also search based on the device name, device model, and device MAC addresses if searching for a Test Agent associated with the device. You can also hover over an interface name to view the details such as the Test Agent name, bind

Table 291: Ping Parameters (Continued)

Parameter	Description
	interface, bind family, and to check if the Test Agent is run on a device or not. You can also remove the interface filters you have selected and reset the filters to the default mode by using the Reset filters button.
Hosts*	Specify the hostname or the destination IP. A host is the remote endpoint to which the Test Agent sends the request. When you click the Host text box, the Select Host page appears where you can enter hostnames. To add more than one hosts, click + Add Host and specify the following: • Host—The hostname or the IP address of the remote endpoint. Maximum length—255 characters • Name—The text box is automatically populated based on the data you specified in the Host text box.
Time between requests*	Specify the time taken between successive ping requests initiated by Test Agent. Unit—Seconds (s) Default value—10.00 s Range—0.01 s to 3600 s
Thresholds for error seconds (ES)	

Table 291: Ping Parameters (Continued)

Parameter	Description
Delay*	Specify the maximum threshold for the delay. If the round-trip time for a Ping response exceeds this value, an errored-second is raised. Unit—Milliseconds (ms) Default value—1000 ms Range—1 ms through 30000 ms
DV (Delay variance)	Specify the maximum threshold value for jitter or delay variance (DV). DV is the difference between the maximum and minimum delay experienced in a one second metrics interval. If the DV exceeds the defined value, an errored-second is raised. Delay variation (DV) occurs when different packets take different amount of time to travel from a Test Agent to a remote endpoint. Packets are sent at regular interval of time and if variation is experienced in consecutive packets, the Test Agent generates an errored-second event. Unit—Milliseconds (ms) Default value—500 ms Range—0 ms through 10000 ms
Advanced	
UDP echo	Enable UDP Echo protocol instead of ICMP Ping. By default, the toggle button is disabled.

Table 291: Ping Parameters *(Continued)*

Parameter	Description
Payload	Specify the size of the ping payload. Payload is the actual data in a request packet. Unit—Bytes Default value—56 bytes Range—0 byte through 65000 bytes
TTL (Time to Live)	Specify the number of times the packets hop before a device discards the packet. Time to live (TTL) value indicates lifespan of a request packet. TTL prevents infinite loop in a network when the packet fails to reach the destination. For example, if you have entered the TTL value as 64, every time the packet passes a device, the value is reduced by one until the packet reaches the remote server. If the entered value reaches zero before reaching the remote endpoint, the packet is discarded. Default value—64 Range—1 through 255
Request Lifetime	Specify the request lifetime value. Request lifetime value is the maximum duration for which a ping request is alive. It determines how long a request persists before it is counted as lost. Unit—Milliseconds (ms) Default value—2000 ms Range—1 ms through 30000 ms

Table 291: Ping Parameters *(Continued)*

Parameter	Description
DSCP/IPP	Specify the Differentiated Services Code Point (DSCP) or the IP Precedence (IPP) value that is used in the IP packet headers. The IPP is the three-bit binary values (Precedence) in the ToS field of the IP header. An IPP value can be in the 0-7 range. IPP value informs the router about the priority of the packet. The higher the IPP value, the more the priority of the packet. See RFC 791 for more information. DSCP is a six-bit binary value in the DS field of the IP header. DSCP value facilitates QoS for traffic management through the Best Effort, Assured Forwarding, Class Selector, and the Expedited Forwarding categories. See RFC 2474 for more information. Default value—0 Range—0 through 255

Table 292: RPM CISCO TWAMP Reflector

Parameter	Description
General	

Table 292: RPM CISCO TWAMP Reflector (Continued)

Parameter	Description
Reflectors*	Select one or more Test Agent interfaces on which you want to run the Test. To select one or more Test Agent interfaces: 1. Click the Clients text box. The Test Agent/Interfaces choices page appears. 2. Select one or more Test Agent interfaces. For information on filtering Test Agent interfaces, see "Search for Test Agent interfaces by using Filters" on page 1136. 3. Click OK. NOTE: On the Test Agent/Interfaces choices page, you can filter the Test Agent interfaces based on the following criteria: • Interface name or the Test Agent name, or both. Based on your requirement, you can select or deselect Test Agents or Interfaces. • IPv4 address (with or without IPv4 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv4 or No IPv4 to view or hide the interfaces that are associated with the IPv4 family and the interfaces that are not associated with the IPv4 family. • IPv6 address (with or without IPv6 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv6 or No IPv6 to view or hide the interfaces that are associated with the IPv6 family and the interfaces that are not associated with the IPv6 family.

Table 292: RPM CISCO TWAMP Reflector *(Continued)*

Parameter	Description
	• Deployment type of a Test Agent: • Application—Filter out the Test Agents deployed in a device as an application. • Device—Filter out the Test Agents associated with a device. Based on your requirement, you can select or deselect Application or Device to view or hide the Test Agents deployed in a device as an application or the Test Agents associated with a device. • Connection types between the interface and the router: • Link—Indicates that the status is Up and there is an active connection for the interface. • No Link—Indicates that the status is Down and there is no active connection for the interface. Based on your requirement, you can select or deselect Link or No Link. • Connection statuses of the Test Agent: • Online—The Test Agent is connected to Routing Director. • Offline—The Test Agent is not connected to Routing Director. When the interface supports RPM plugins, the interface contains <i>(On device)</i> at the end of its name indicating the Device Online status.

Table 292: RPM CISCO TWAMP Reflector (Continued)

Parameter	Description
	When the interface supports only the normal plugins, the interface will not display (<i>On device</i>) at the end of its name, and the status refers to the Online status. Based on your requirement, you can select or deselect Online or Offline. - Interface types of the Test Agent: - Normal—Network interfaces that are used for testing traffic. - Management—Network interfaces used by the Test Agents to connect to Routing Director. - Unknown—Network interfaces that the Test Agent is unable to identify or that have been removed. Based on your requirement, you can select or deselect Normal, Management, or Unknown. You can enable the Display Device Names toggle button to display the hostnames of the routers that are associated with the Test Agents. You can also enable the Hide unsupported toggle button to hide devices that are not supported by the plugin you have selected. To search for a specific interface or a Test Agent, enter one or more keywords in the search text box. You can search based on the name, description, IP and MAC address of the interfaces or the Test Agents. You can also search based on the device name, device model, and device MAC addresses if searching for a Test Agent associated with the device. You can also hover over an interface name to view the details such as the Test Agent name, bind

Table 292: RPM CISCO TWAMP Reflector *(Continued)*

Parameter	Description
	interface, bind family, and to check if the Test Agent is run on a device or not. You can also remove the interface filters you have selected and reset the filters to the default mode by using the Reset filters button.
Ports*	Specify the Test session port value. When you click the text box, the Select Test Session Port Value page appears where you can add port values. On this page: • Test Session Port Value—Specify the local port used for the test session. Maximum length—65535 characters Minimum length—1 Default length—862 • Host—Specify IPv4 or IPv6 Address of TWAMP client. If no address is specified, Routing Director accepts connections from any TWAMP sender or client. Maximum length—255 characters • Client Test session port—Specify the port used by the client for the test session. If no port value is specified, Routing Director accepts connections from any port of the TWAMP client. Maximum length—65535 characters Minimum length—1 • Name—This text box is automatically populated based on the data you specified in the Host text box.
Advanced	

Table 292: RPM CISCO TWAMP Reflector (Continued)

Parameter	Description
Collection interval*	Specify the collection interval. Collection interval is the frequency at which the results are collected from the device. Collection interval must be larger than the value specified in the Timeout period field. Unit—Seconds (s) Default value—30s Range—30s through 300s
Session Timeout	Specify the duration of time after which an idle session is automatically ended. Unit—Seconds (s) Default value—35s Range— 10s through 300s
Timeout period	Specify the inactivity timeout period for the test session. The session terminates if no activity happens within the specified time. Unit—Seconds (s) Range—1s through 604800s
Inactivity-timer value	Specify the inactivity timeout period for the TWAMP server. The session terminates if no activity happens within the specified time. Unit—Seconds (s) Range—1s through 6000s
Server Test Session Port IP	Specify the IPv4 or IPv6 address of the TWAMP reflector used for test sessions. Maximum length—255 characters

Table 292: RPM CISCO TWAMP Reflector *(Continued)*

Parameter	Description
Control Session Port	Specify the port used for the control session. Specify 0 for TWAMP Light and specify any value starting from 862 through 65535 for managed TWAMP. Default value—0 Range—0 through 65535
Session Duration	Specify the total duration the test session runs. Unit—Seconds (s) Default value—Unlimited (the session continues until the session is manually stopped.) Range— 60s through 86400s
VRF	Specify the VRF (Virtual Routing and Forwarding) instance to use for the test session. Default value—Any Maximum Length—255 characters
Session ID	Specify the session ID used for the TWAMP light test session. Default value—1 Range—1 through 65535

Table 293: IPTV MPEG Parameters

Parameter	Description
General	

Table 293: IPTV MPEG Parameters (Continued)

Parameter	Description
Clients*	Select one or more Test Agent interfaces on which you want to run the Monitor. To select one or more Test Agent interfaces: 1. Click the Clients text box. The Test Agent/Interfaces choices page appears. 2. Select one or more Test Agent interfaces. For information on filtering Test Agent interfaces, see "Search for Test Agent interfaces by using Filters" on page 1142. 3. Click OK. NOTE: On the Test Agent/Interfaces choices page, you can filter the Test Agent interfaces based on the following criteria: • Interface name or the Test Agent name, or both. Based on your requirement, you can select or deselect Test Agents or Interfaces. • IPv4 address (with or without IPv4 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv4 or No IPv4 to view or hide the interfaces that are associated with the IPv4 family and the interfaces that are not associated with the IPv4 family. • IPv6 address (with or without IPv6 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv6 or No IPv6 to view or hide the interfaces that are associated with the IPv6 family and the interfaces that are not associated with the IPv6 family.

Table 293: IPTV MPEG Parameters *(Continued)*

Parameter	Description
	• Deployment type of a Test Agent: • Application—Filter out the Test Agents deployed in a device as an application. • Device—Filter out the Test Agents associated with a device. Based on your requirement, you can select or deselect Application or Device to view or hide the Test Agents deployed in a device as an application or the Test Agents associated with a device. • Connection types between the interface and the router: • Link—Indicates that the status is Up and there is an active connection for the interface. • No Link—Indicates that the status is Down and there is no active connection for the interface. Based on your requirement, you can select or deselect Link or No Link. • Connection statuses of the Test Agent: • Online—The Test Agent is connected to Routing Director. • Offline—The Test Agent is not connected to Routing Director. When the interface supports RPM plugins, the interface contains <i>(On device)</i> at the end of its name indicating the Device Online status.

Table 293: IPTV MPEG Parameters *(Continued)*

Parameter	Description
	When the interface supports only the normal plugins, the interface will not display <i>(On device)</i> at the end of its name, and the status refers to the Online status. Based on your requirement, you can select or deselect Online or Offline. - Interface types of the Test Agent: - Normal—Network interfaces that are used for testing traffic. - Management—Network interfaces used by the Test Agents to connect to Routing Director. - Unknown—Network interfaces that the Test Agent is unable to identify or that have been removed. Based on your requirement, you can select or deselect Normal, Management, or Unknown. You can enable the Display Device Names toggle button to display the hostnames of the routers that are associated with the Test Agents. You can also enable the Hide unsupported toggle button to hide devices that are not supported by the plugin you have selected. To search for a specific interface or a Test Agent, enter one or more keywords in the search text box. You can search based on the name, description, IP and MAC address of the interfaces or the Test Agents. You can also search based on the device name, device model, and device MAC addresses if searching for a Test Agent associated with the device. You can also hover over an interface name to view the details such as the Test Agent name, bind

Table 293: IPTV MPEG Parameters (Continued)

Parameter	Description
	interface, bind family, and to check if the Test Agent is run on a device or not. You can also remove the interface filters you have selected and reset the filters to the default mode by using the Reset filters button.
IPTV Channels*	Specify one or more IPTV channels to monitor from the preset IPTV channel list. When you click IPTV Channels, the Select Channel page appears where you can enter channels. To add more than one channels, click + Add channel and specify the following: • Multicast group address—Specify the IPv4 or IPv6 multicast address of the IPTV channel. Maximum length—64 characters. • Multicast group port—Specify the UDP destination port used in the IPTV multicast stream. Range—1 through 65535. • Multicast group source—Specify the source address of the multicast group. When provided, IGMPv3 (for IPv4) or MLDv2 (for IPv6) is used for joining the channels. • Program number—Specify the program number of the channel within a multicast group. • Name—Specify the name of the IPTV channel.
Thresholds for error seconds (ES)	

Table 293: IPTV MPEG Parameters *(Continued)*

Parameter	Description
MPEG loss (CC errors/s)	Specify the maximum allowed MPEG packet loss (Continuity Count errors) per second. An errored-second is raised if the MPEG loss exceeds this threshold. Continuity Counter (CC) error indicates an issue with the packet sequence in an MPEG transport stream.
PAT/PMT delay variation (ms)	Specify the maximum tolerated jitter (delay variation) of the received MPEG stream. An errored-second is triggered if the PCR jitter or RTP jitter exceeds this threshold. PCR provides timestamps to maintain synchronization of audio and video, RTP provides end-to-end delivery of real-time data over IP networks.
PAT/PMT interval (s)	Specify the maximum allowable time interval between Program Allocation Table (PAT)/ Program Map Table (PMT) transmissions. PAT has information on all the available programs in a transport stream. PAT maps all the programs numbers to PMT packet identifiers. PMT has information on the streams and lists the PIDs of all the streams associated with a program.
PID interval (s)	Specify the maximum allowable time interval between audio or video Packet Identifiers (PIDs) as specified by Program Map Table (PMT). PID is a 13-bit field in the header of a transport stream used to identify different data streams, such as video, audio, subtitles, within a transport stream.

Table 294: Netflix Speedtest Parameters

Parameter	Description
General	

Table 294: Netflix Speedtest Parameters (Continued)

Parameter	Description
Clients*	Select one or more Test Agent interfaces on which you want to run the Monitor. To select one or more Test Agent interfaces: 1. Click the Clients text box. The Test Agent/Interfaces choices page appears. 2. Select one or more Test Agent interfaces. For information on filtering Test Agent interfaces, see "Search for Test Agent interfaces by using Filters." on page 1148 3. Click OK. NOTE: On the Test Agent/Interfaces choices page, you can filter the Test Agent interfaces based on the following criteria: • Interface name or the Test Agent name, or both. Based on your requirement, you can select or deselect Test Agents or Interfaces. • IPv4 address (with or without IPv4 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv4 or No IPv4 to view or hide the interfaces that are associated with the IPv4 family and the interfaces that are not associated with the IPv4 family. • IPv6 address (with or without IPv6 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv6 or No IPv6 to view or hide the interfaces that are associated with the IPv6 family and the interfaces that are not associated with the IPv6 family.

Table 294: Netflix Speedtest Parameters *(Continued)*

Parameter	Description
	• Deployment type of a Test Agent: • Application—Filter out the Test Agents deployed in a device as an application. • Device—Filter out the Test Agents associated with a device. Based on your requirement, you can select or deselect Application or Device to view or hide the Test Agents deployed in a device as an application or the Test Agents associated with a device. • Connection types between the interface and the router: • Link—Indicates that the status is Up and there is an active connection for the interface. • No Link—Indicates that the status is Down and there is no active connection for the interface. Based on your requirement, you can select or deselect Link or No Link. • Connection statuses of the Test Agent: • Online—The Test Agent is connected to Routing Director. • Offline—The Test Agent is not connected to Routing Director. When the interface supports RPM plugins, the interface contains <i>(On device)</i> at the end of its name indicating the Device Online status.

Table 294: Netflix Speedtest Parameters (Continued)

Parameter	Description
	When the interface supports only the normal plugins, the interface will not display (<i>On device</i>) at the end of its name, and the status refers to the Online status. Based on your requirement, you can select or deselect Online or Offline. - Interface types of the Test Agent: - Normal—Network interfaces that are used for testing traffic. - Management—Network interfaces used by the Test Agents to connect to Routing Director. - Unknown—Network interfaces that the Test Agent is unable to identify or that have been removed. Based on your requirement, you can select or deselect Normal, Management, or Unknown. You can enable the Display Device Names toggle button to display the hostnames of the routers that are associated with the Test Agents. You can also enable the Hide unsupported toggle button to hide devices that are not supported by the plugin you have selected. To search for a specific interface or a Test Agent, enter one or more keywords in the search text box. You can search based on the name, description, IP and MAC address of the interfaces or the Test Agents. You can also search based on the device name, device model, and device MAC addresses if searching for a Test Agent associated with the device. You can also hover over an interface name to view the details such as the Test Agent name, bind

Table 294: Netflix Speedtest Parameters *(Continued)*

Parameter	Description
	interface, bind family, and to check if the Test Agent is run on a device or not. You can also remove the interface filters you have selected and reset the filters to the default mode by using the Reset filters button.
Measurement Cycle period	Specify the time interval between successive measurement cycles. On each measurement cycle, the Test Agent downloads Netflix test segments over HTTPS from one or several OCAs (Open Connect Appliances) operated by Netflix. OCA is a server appliance to cache and deliver content locally to subscribers, reducing the load on Netflix's central servers. Unit—Minutes (min) Default value—15 min Range—1 min through 60 min
Upload or Download	Select the type of bandwidth to measure the rate of the data transfer. Upload data transfer is calculated from the Test Agent to OCAs, while download is calculated from the OCAs to the Test Agent.
Advanced	
Minimum Duration (s)	Specify the minimum duration of speedtest that must be run before checking for throughput stabilization. Unit—Seconds (s) Default value—2 s Range—1 s through 3600 s

Table 294: Netflix Speedtest Parameters *(Continued)*

Parameter	Description
Maximum Duration (s)	Specify the maximum duration of speedtest that must be run to check for throughput stabilization. Unit—Seconds (s) Default value—30 s Range—5 s through 3600 s
Minimum Concurrent OCA connections	Specify the minimum number of concurrent Open Connect Appliance (OCA) connections used for bandwidth testing. Default value—3. Range—1 through 5.
Maximum Concurrent OCA connections	Specify the maximum number of concurrent Open Connect Appliance (OCA) connections used for bandwidth testing. Default value—5 Range—1 through 5
Thresholds for error seconds (ES)	
Minimum bandwidth (Mbit/s)	Specify the minimum expected bandwidth during the speedtest. An errored-second is raised if the bandwidth drops below the defined value.
Maximum latency (ms)	Specify the minimum expected bandwidth during the speedtest. An errored-second is raised if the bandwidth drops below the defined value.

Table 295: OTT - HLS Parameters

Parameter	Description
General	

Table 295: OTT - HLS Parameters (Continued)

Parameter	Description
Clients*	Select one or more Test Agent interfaces on which you want to run the Monitor. To select one or more Test Agent interfaces: 1. Click the Clients text box. The Test Agent/Interfaces choices page appears. 2. Select one or more Test Agent interfaces. For information on filtering Test Agent interfaces, see "Search for Test Agent interfaces by using Filters" on page 1154. 3. Click OK. NOTE: On the Test Agent/Interfaces choices page, you can filter the Test Agent interfaces based on the following criteria: • Interface name or the Test Agent name, or both. Based on your requirement, you can select or deselect Test Agents or Interfaces. • IPv4 address (with or without IPv4 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv4 or No IPv4 to view or hide the interfaces that are associated with the IPv4 family and the interfaces that are not associated with the IPv4 family. • IPv6 address (with or without IPv6 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv6 or No IPv6 to view or hide the interfaces that are associated with the IPv6 family and the interfaces that are not associated with the IPv6 family.

Table 295: OTT - HLS Parameters *(Continued)*

Parameter	Description
	• Deployment type of a Test Agent: • Application—Filter out the Test Agents deployed in a device as an application. • Device—Filter out the Test Agents associated with a device. Based on your requirement, you can select or deselect Application or Device to view or hide the Test Agents deployed in a device as an application or the Test Agents associated with a device. • Connection types between the interface and the router: • Link—Indicates that the status is Up and there is an active connection for the interface. • No Link—Indicates that the status is Down and there is no active connection for the interface. Based on your requirement, you can select or deselect Link or No Link. • Connection statuses of the Test Agent: • Online—The Test Agent is connected to Routing Director. • Offline—The Test Agent is not connected to Routing Director. When the interface supports RPM plugins, the interface contains <i>(On device)</i> at the end of its name indicating the Device Online status.

Table 295: OTT - HLS Parameters (Continued)

Parameter	Description
	When the interface supports only the normal plugins, the interface will not display (<i>On device</i>) at the end of its name, and the status refers to the Online status. Based on your requirement, you can select or deselect Online or Offline. - Interface types of the Test Agent: - Normal—Network interfaces that are used for testing traffic. - Management—Network interfaces used by the Test Agents to connect to Routing Director. - Unknown—Network interfaces that the Test Agent is unable to identify or that have been removed. Based on your requirement, you can select or deselect Normal, Management, or Unknown. You can enable the Display Device Names toggle button to display the hostnames of the routers that are associated with the Test Agents. You can also enable the Hide unsupported toggle button to hide devices that are not supported by the plugin you have selected. To search for a specific interface or a Test Agent, enter one or more keywords in the search text box. You can search based on the name, description, IP and MAC address of the interfaces or the Test Agents. You can also search based on the device name, device model, and device MAC addresses if searching for a Test Agent associated with the device. You can also hover over an interface name to view the details such as the Test Agent name, bind

Table 295: OTT - HLS Parameters (Continued)

Parameter	Description
	interface, bind family, and to check if the Test Agent is run on a device or not. You can also remove the interface filters you have selected and reset the filters to the default mode by using the Reset filters button.
URLs*	Specify the URL of the video stream to which you send the connection requests. The URL can be either a playlist linking to other playlists for the different variants, or a playlist containing the segments. The file extension is .m3u or .m3u8 When you click URLs, the Select URL page appears where you can enter URLs. To add more than one URLs, click + Add URL and specify the following: • URL—The host or the IP address of the remote endpoint. • Name—The text box is automatically populated based on the data you specified in URL.
Buffer size (s)	Specify the target duration of buffered data. If the duration of the buffer data drops below the defined value, new segments are downloaded. Unit—Seconds (ms) Default value—60 s Minimum—1 s

Table 295: OTT - HLS Parameters (*Continued*)

Parameter	Description
Initial Buffering (s)	Specify the duration of the initial buffered data required before the playback begins. Unit—Seconds (ms) Default value—10 s Minimum—1 s
Loop	Enable or disable the Loop button. If enabled, the playback loops when the stream reaches the end of the playlist. If disabled, the playback stops, and errored-seconds are raised. By default, the toggle button is enabled.
Thresholds for error seconds (ES)	
Minimum playback rate (Mbit/s)	Specify the minimum expected playback rate of the video stream. An errored second is triggered if its value drops below this threshold. Unit—Megabit per second (Mbit/s) Minimum value—0 Mbit/s
Minimum download rate (Mbit/s)	Specify the minimum expected playback rate of the video stream. An errored-second is raised if the download rate drops below the defined value. Unit—Megabit per second (Mbit/s) Minimum value—0 Mbit/s

Table 295: OTT - HLS Parameters *(Continued)*

Parameter	Description
Minimum selected rate (Mbit/s)	Specify the minimum expected download rate for the segments. An errored-second is raised if the selected rate drops below the defined value Variants are the versions of media streams that is encoded at a specific rate. They enable adaptive streaming where the Test Agents can select variants to provide the best experience. Unit—Megabit per second (Mbit/s) Minimum value—0 Mbit/s
Minimum buffer size (s)	Specify the minimum expected duration of the data segment buffered in the Test Agent. An errored-second is triggered if its value drops below the defined value. Unit—Megabit per second (Mbit/s) Minimum value—0 Mbit/s

Table 296: TCP Parameters

Parameter	Description
General	

Table 296: TCP Parameters (Continued)

Parameter	Description
Clients*	Select one or more Test Agent interfaces which you want to use as TCP clients. To select one or more Test Agent interfaces: 1. Click Clients. The Test Agent/Interfaces choices page appears. 2. Select one or more Test Agent interfaces. For information on filtering Test Agent interfaces, see "Search for Test Agent interfaces by using Filters" on page 1160. 3. Click OK. NOTE: On the Test Agent/Interfaces choices page, you can filter the Test Agent interfaces based on the following criteria: • Interface name or the Test Agent name, or both. Based on your requirement, you can select or deselect Test Agents or Interfaces. • IPv4 address (with or without IPv4 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv4 or No IPv4 to view or hide the interfaces that are associated with the IPv4 family and the interfaces that are not associated with the IPv4 family. • IPv6 address (with or without IPv6 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv6 or No IPv6 to view or hide the interfaces that are associated with the IPv6 family and the interfaces that are not associated with the IPv6 family.

Table 296: TCP Parameters *(Continued)*

Parameter	Description
	• Deployment type of a Test Agent: • Application—Filter out the Test Agents deployed in a device as an application. • Device—Filter out the Test Agents associated with a device. Based on your requirement, you can select or deselect Application or Device to view or hide the Test Agents deployed in a device as an application or the Test Agents associated with a device. • Connection types between the interface and the router: • Link—Indicates that the status is Up and there is an active connection for the interface. • No Link—Indicates that the status is Down and there is no active connection for the interface. Based on your requirement, you can select or deselect Link or No Link. • Connection statuses of the Test Agent: • Online—The Test Agent is connected to Routing Director. • Offline—The Test Agent is not connected to Routing Director. When the interface supports RPM plugins, the interface contains <i>(On device)</i> at the end of its name indicating the Device Online status.

Table 296: TCP Parameters *(Continued)*

Parameter	Description
	When the interface supports only the normal plugins, the interface will not display <i>(On device)</i> at the end of its name, and the status refers to the Online status. Based on your requirement, you can select or deselect Online or Offline. - Interface types of the Test Agent: - Normal—Network interfaces that are used for testing traffic. - Management—Network interfaces used by the Test Agents to connect to Routing Director. - Unknown—Network interfaces that the Test Agent is unable to identify or that have been removed. Based on your requirement, you can select or deselect Normal, Management, or Unknown. You can enable the Display Device Names toggle button to display the hostnames of the routers that are associated with the Test Agents. You can also enable the Hide unsupported toggle button to hide devices that are not supported by the plugin you have selected. To search for a specific interface or a Test Agent, enter one or more keywords in the search text box. You can search based on the name, description, IP and MAC address of the interfaces or the Test Agents. You can also search based on the device name, device model, and device MAC addresses if searching for a Test Agent associated with the device. You can also hover over an interface name to view the details such as the Test Agent name, bind

Table 296: TCP Parameters *(Continued)*

Parameter	Description
	interface, bind family, and to check if the Test Agent is run on a device or not. You can also remove the interface filters you have selected and reset the filters to the default mode by using the Reset filters button.

Table 296: TCP Parameters (Continued)

Parameter	Description
Servers*	Select one or more Test Agent interfaces that you want to use as TCP servers. To select one or more Test Agent interfaces: 1. Click Clients. The Test Agent/Interfaces choices page appears. 2. Select one or more Test Agent interfaces. For information on filtering Test Agents, see "Search for Test Agent interfaces by using Filters" on page 1164. 3. Click OK. NOTE: On the Test Agent/Interfaces choices page, you can filter the Test Agent interfaces based on the following criteria: • Interface name or the Test Agent name, or both. Based on your requirement, you can select or deselect Test Agents or Interfaces. • IPv4 address (with or without IPv4 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv4 or No IPv4 to view or hide the interfaces that are associated with the IPv4 family and the interfaces that are not associated with the IPv4 family. • IPv6 address (with or without IPv6 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv6 or No IPv6 to view or hide the interfaces that are associated with the IPv6 family and the interfaces that are not associated with the IPv6 family.

Table 296: TCP Parameters (Continued)

Parameter	Description
	• Deployment type of a Test Agent: • Application—Filter out the Test Agents deployed in a device as an application. • Device—Filter out the Test Agents associated with a device. Based on your requirement, you can select or deselect Application or Device to view or hide the Test Agents deployed in a device as an application or the Test Agents associated with a device. • Connection types between the interface and the router: • Link—Indicates that the status is Up and there is an active connection for the interface. • No Link—Indicates that the status is Down and there is no active connection for the interface. Based on your requirement, you can select or deselect Link or No Link. • Connection statuses of the Test Agent: • Online—The Test Agent is connected to Routing Director. • Offline—The Test Agent is not connected to Routing Director. When the interface supports RPM plugins, the interface contains (<i>On device</i>) at the end of its name indicating the Device Online status.

Table 296: TCP Parameters *(Continued)*

Parameter	Description
	When the interface supports only the normal plugins, the interface will not display <i>(On device)</i> at the end of its name, and the status refers to the Online status. Based on your requirement, you can select or deselect Online or Offline. - Interface types of the Test Agent: - Normal—Network interfaces that are used for testing traffic. - Management—Network interfaces used by the Test Agents to connect to Routing Director. - Unknown—Network interfaces that the Test Agent is unable to identify or that have been removed. Based on your requirement, you can select or deselect Normal, Management, or Unknown. You can enable the Display Device Names toggle button to display the hostnames of the routers that are associated with the Test Agents. You can also enable the Hide unsupported toggle button to hide devices that are not supported by the plugin you have selected. To search for a specific interface or a Test Agent, enter one or more keywords in the search text box. You can search based on the name, description, IP and MAC address of the interfaces or the Test Agents. You can also search based on the device name, device model, and device MAC addresses if searching for a Test Agent associated with the device. You can also hover over an interface name to view the details such as the Test Agent name, bind

Table 296: TCP Parameters *(Continued)*

Parameter	Description
	interface, bind family, and to check if the Test Agent is run on a device or not. You can also remove the interface filters you have selected and reset the filters to the default mode by using the Reset filters button.
Direction	Choose the direction of the test traffic. The data is sent from 'Server to Client' or 'Client to Server'. Client to server—evaluates the data upload capabilities. Server to client—evaluates the data download capabilities.
Rate (Mbit/s)	Rate indicates the target data rate between the server and the client Test Agents. If you do not define a value, the rate will not be limited, and the TCP uses TCP congestion control. Unit—Megabit per second (Mbit/s) Range—0.01 Mbit/s through 1000000.0 Mbit/s
Advanced	
Server TCP port	Specify the port number used by the server Test Agent. Default value—5000 Range—1 through 65535
Client TCP port	Specify the port number used by the client Test Agent. Range—1 through 65535

Table 296: TCP Parameters *(Continued)*

Parameter	Description
Number of flows	Specify the total number of TCP sessions. If more than one flows are specified, you should not set the Client port. If set, the plug-in activates an invalid configuration error, and if not set, the Client port gets a randomly assigned port number. Default value—1 Range—1 through 64
DSCP/IPP	Specify the Differentiated Services Code Point (DSCP) or the IP Precedence (IPP) value that is used in the IP packet headers. DSCP is a six-bit binary value in the DS field of the IP header. DSCP value facilitates QoS for traffic management through the Best Effort, Assured Forwarding, Class Selector, and the Expedited Forwarding categories. See RFC 2474 for more information. Default value—0 Range—0 through 63
VLAN priority (PCP)	Specify the value for the VLAN priority code point to be used in VLAN header. Range—0 through 7
Client connect delay (s)	Specify the allowable delay until the client initiates TCP connection towards the server. Unit— Seconds (s) Default value—0 Range—0 through 10
Proxy address	Specify the IP address of the TCP proxy server.

Table 296: TCP Parameters *(Continued)*

Parameter	Description
Proxy port	Specify the port number that the Test Agent uses for TCP proxy server. Range—1 through 65535
Thresholds for error seconds (ES)	
Min rate (Mbit/s)	Specify the minimum expected rate. An errored-second is raised if the TCP data rate drops below the defined value. Unit—Megabit per second (Mbit/s) Minimum—0.001 Mbit/s
Max rate (Mbit/s)	Specify the minimum expected rate. An errored-second is raised if the TCP data rate drops below the defined value. Unit—Megabit per second (Mbit/s) Minimum—0.001 Mbit/s

Table 297: UDP Parameters

Parameter	Description
General	

Table 297: UDP Parameters *(Continued)*

Parameter	Description
Clients*	Select one or more Test Agent interfaces which you want to use as UDP clients. To select one or more Test Agent interfaces: 1. Click Clients. The Test Agent/Interfaces choices page appears. 2. Select one or more Test Agent interfaces. For information on filtering Test Agent interfaces, see "Search for Test Agent interfaces by using Filters" on page 1170. 3. Click OK. NOTE: On the Test Agent/Interfaces choices page, you can filter the Test Agent interfaces based on the following criteria: • Interface name or the Test Agent name, or both. Based on your requirement, you can select or deselect Test Agents or Interfaces. • IPv4 address (with or without IPv4 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv4 or No IPv4 to view or hide the interfaces that are associated with the IPv4 family and the interfaces that are not associated with the IPv4 family. • IPv6 address (with or without IPv6 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv6 or No IPv6 to view or hide the interfaces that are associated with the IPv6 family and the interfaces that are not associated with the IPv6 family.

Table 297: UDP Parameters *(Continued)*

Parameter	Description
	• Deployment type of a Test Agent: • Application—Filter out the Test Agents deployed in a device as an application. • Device—Filter out the Test Agents associated with a device. Based on your requirement, you can select or deselect Application or Device to view or hide the Test Agents deployed in a device as an application or the Test Agents associated with a device. • Connection types between the interface and the router: • Link—Indicates that the status is Up and there is an active connection for the interface. • No Link—Indicates that the status is Down and there is no active connection for the interface. Based on your requirement, you can select or deselect Link or No Link. • Connection statuses of the Test Agent: • Online—The Test Agent is connected to Routing Director. • Offline—The Test Agent is not connected to Routing Director. When the interface supports RPM plugins, the interface contains <i>(On device)</i> at the end of its name indicating the Device Online status.

Table 297: UDP Parameters *(Continued)*

Parameter	Description
	When the interface supports only the normal plugins, the interface will not display <i>(On device)</i> at the end of its name, and the status refers to the Online status. Based on your requirement, you can select or deselect Online or Offline. - Interface types of the Test Agent: - Normal—Network interfaces that are used for testing traffic. - Management—Network interfaces used by the Test Agents to connect to Routing Director. - Unknown—Network interfaces that the Test Agent is unable to identify or that have been removed. Based on your requirement, you can select or deselect Normal, Management, or Unknown. You can enable the Display Device Names toggle button to display the hostnames of the routers that are associated with the Test Agents. You can also enable the Hide unsupported toggle button to hide devices that are not supported by the plugin you have selected. To search for a specific interface or a Test Agent, enter one or more keywords in the search text box. You can search based on the name, description, IP and MAC address of the interfaces or the Test Agents. You can also search based on the device name, device model, and device MAC addresses if searching for a Test Agent associated with the device. You can also hover over an interface name to view the details such as the Test Agent name, bind

Table 297: UDP Parameters *(Continued)*

Parameter	Description
	interface, bind family, and to check if the Test Agent is run on a device or not. You can also remove the interface filters you have selected and reset the filters to the default mode by using the Reset filters button.

Table 297: UDP Parameters *(Continued)*

Parameter	Description
Server*	Select one or more Test Agent interfaces which you want to use as UDP servers. To select one or more Test Agent interfaces: 1. Click Clients. The Test Agent/Interfaces choices page appears. 2. Select one or more Test Agent interfaces. For information on filtering Test Agent interfaces, see "Search for Test Agent interfaces by using Filters" on page 1174. 3. Click OK. NOTE: On the Test Agent/Interfaces choices page, you can filter the Test Agent interfaces based on the following criteria: • Interface name or the Test Agent name, or both. Based on your requirement, you can select or deselect Test Agents or Interfaces. • IPv4 address (with or without IPv4 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv4 or No IPv4 to view or hide the interfaces that are associated with the IPv4 family and the interfaces that are not associated with the IPv4 family. • IPv6 address (with or without IPv6 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv6 or No IPv6 to view or hide the interfaces that are associated with the IPv6 family and the interfaces that are not associated with the IPv6 family.

Table 297: UDP Parameters *(Continued)*

Parameter	Description
	- Deployment type of a Test Agent: - Application—Filter out the Test Agents deployed in a device as an application. - Device—Filter out the Test Agents associated with a device. Based on your requirement, you can select or deselect Application or Device to view or hide the Test Agents deployed in a device as an application or the Test Agents associated with a device. - Connection types between the interface and the router: - Link—Indicates that the status is Up and there is an active connection for the interface. - No Link—Indicates that the status is Down and there is no active connection for the interface. Based on your requirement, you can select or deselect Link or No Link. - Connection statuses of the Test Agent: - Online—The Test Agent is connected to Routing Director. - Offline—The Test Agent is not connected to Routing Director. When the interface supports RPM plugins, the interface contains <i>(On device)</i> at the end of its name indicating the Device Online status.

Table 297: UDP Parameters *(Continued)*

Parameter	Description
	When the interface supports only the normal plugins, the interface will not display <i>(On device)</i> at the end of its name, and the status refers to the Online status. Based on your requirement, you can select or deselect Online or Offline. - Interface types of the Test Agent: - Normal—Network interfaces that are used for testing traffic. - Management—Network interfaces used by the Test Agents to connect to Routing Director. - Unknown—Network interfaces that the Test Agent is unable to identify or that have been removed. Based on your requirement, you can select or deselect Normal, Management, or Unknown. You can enable the Display Device Names toggle button to display the hostnames of the routers that are associated with the Test Agents. You can also enable the Hide unsupported toggle button to hide devices that are not supported by the plugin you have selected. To search for a specific interface or a Test Agent, enter one or more keywords in the search text box. You can search based on the name, description, IP and MAC address of the interfaces or the Test Agents. You can also search based on the device name, device model, and device MAC addresses if searching for a Test Agent associated with the device. You can also hover over an interface name to view the details such as the Test Agent name, bind

Table 297: UDP Parameters *(Continued)*

Parameter	Description
	interface, bind family, and to check if the Test Agent is run on a device or not. You can also remove the interface filters you have selected and reset the filters to the default mode by using the Reset filters button.
Direction	Choose the direction of the test traffic. The data is sent from 'Server to Client' or 'Client to Server'. Client to server—Evaluates the data upload capabilities. Server to client—Evaluates the data download capabilities.
Rate (Mbit/s)	Specify the target data rate at which the client or server Test Agent transmits the UDP frames. Unit—Megabit per second (Mbit/s) Range—0.01 Mbit/s through 1000000.0 Mbit/s
Advanced	
Server UDP port	Specify the port number of the UDP server to which the client Test Agent form a connection. Default value—5000 Range—1 through 65535
Client UDP port	Specify the port number of the UDP server to which the client Test Agent form a connection. Range—1 through 65535

Table 297: UDP Parameters *(Continued)*

Parameter	Description
Number of flows	Specify the number of UDP flows. If more than one flows are specified, you should not set the Client port. If set, the plug-in activates an invalid configuration error, and if not set, the Client port gets a randomly assigned port number. Default value—1 Range—1 through 64
Don't Fragment flag	Enable the Don't Fragment Flag (DF Flag) to restrict the fragmentation of the packets that exceed the MTU. DF Flag is configured in an IP header. Router drops the packet if fragmentation is needed. Enabling the toggle button may cause performance degradation both in the network and in the sending or receiving Test Agents. By default, the toggle button is disabled.
Ethernet frame size (Bytes)	Specify the size of Layer 2 Ethernet frame for the data flow. Frame Size indicates the total size of the data frame sent from a remote endpoint to the Test agent as a response to a request. Default—1518 Range—64 bytes through 9018 bytes

Table 297: UDP Parameters *(Continued)*

Parameter	Description
DSCP/IPP	Specify the Differentiated Services Code Point (DSCP) or the IP Precedence (IPP) value that is used in the IP packet headers. DSCP is a six-bit binary value in the DS field of the IP header. DSCP value facilitates QoS for traffic management through the Best Effort, Assured Forwarding, Class Selector, and the Expedited Forwarding categories. See RFC 2474 for more information. Default value—0 Range—0 through 63
VLAN priority (PCP)	Specify the value for the VLAN priority code point to be used in VLAN header. Range—0 through 7
Thresholds for error seconds (ES)	
Loss (%)	Specify the loss percentage value. If the loss percentage exceeds the configured value, an errored-second is triggered. Unit—Percentage (%) Range—0.0% through 100.0%
Delay (ms)	Specify the threshold value for the minimum delay measured between the Test Agent client and server. If delay exceeds the defined value, an errored-seconds is raised. Unit—Milliseconds (ms) Minimum—0.0 ms

Table 297: UDP Parameters *(Continued)*

Parameter	Description
Jitter (ms)	Specify the maximum threshold value for delay variance (jitter) in UDP transaction. Delay variation (DV) occurs when different packets take different amount of time to travel from a Test Agent to a remote endpoint. Unit—Milliseconds (ms) Minimum—0.0 ms
Expected DSCP	Specify the expected DSCP that must be in the frame when received by Test Agent. If the received DSCP value does not match the configured value, an errored-second is raised. Range—0 through 64
Thresholds for severely error seconds (SES)	
Loss (%)	Specify the maximum allowable loss percentage value. If the loss percentage exceeds the configured value, an errored-second is raised. Unit—Percentage (%) Range—0.0% through 100.0%
Delay (ms)	Specify the maximum allowable value for delay measurement. If the delay between server and client exceeds the configured value during a one-second interval, a severely-error seconds is indicated. Unit—Milliseconds (ms) Minimum—0.0 ms

Table 297: UDP Parameters *(Continued)*

Parameter	Description
Jitter (ms)	Specify the maximum threshold value for delay variance (jitter). Packets are sent at regular interval of time and if variation is experienced in consecutive packets, the Test Agent generates an errored-second event. Unit—Milliseconds (ms) Minimum—0.0 ms

Table 298: TWAMP/TWAMP Light Parameters

Parameter	Description
General	

Table 298: TWAMP/TWAMP Light Parameters (Continued)

Parameter	Description
Senders*	Select one or more Test Agent interfaces on which you want to run the Test. To select one or more Test Agent interfaces: 1. Click the Clients text box. The Test Agent/Interfaces choices page appears. 2. Select one or more Test Agent interfaces. For information on filtering Test Agent interfaces, see "Search for Test Agent interfaces by using Filters" on page 1182. 3. Click OK. NOTE: On the Test Agent/Interfaces choices page, you can filter the Test Agent interfaces based on the following criteria: • Interface name or the Test Agent name, or both. Based on your requirement, you can select or deselect Test Agents or Interfaces. • IPv4 address (with or without IPv4 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv4 or No IPv4 to view or hide the interfaces that are associated with the IPv4 family and the interfaces that are not associated with the IPv4 family. • IPv6 address (with or without IPv6 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv6 or No IPv6 to view or hide the interfaces that are associated with the IPv6 family and the interfaces that are not associated with the IPv6 family.

Table 298: TWAMP/TWAMP Light Parameters (Continued)

Parameter	Description
	• Deployment type of a Test Agent: • Application—Filter out the Test Agents deployed in a device as an application. • Device—Filter out the Test Agents associated with a device. Based on your requirement, you can select or deselect Application or Device to view or hide the Test Agents deployed in a device as an application or the Test Agents associated with a device. • Connection types between the interface and the router: • Link—Indicates that the status is Up and there is an active connection for the interface. • No Link—Indicates that the status is Down and there is no active connection for the interface. Based on your requirement, you can select or deselect Link or No Link. • Connection statuses of the Test Agent: • Online—The Test Agent is connected to Routing Director. • Offline—The Test Agent is not connected to Routing Director. When the interface supports RPM plugins, the interface contains (<i>On device</i>) at the end of its name indicating the Device Online status.

Table 298: TWAMP/TWAMP Light Parameters (Continued)

Parameter	Description
	When the interface supports only the normal plugins, the interface will not display (<i>On device</i>) at the end of its name, and the status refers to the Online status. Based on your requirement, you can select or deselect Online or Offline. - Interface types of the Test Agent: - Normal—Network interfaces that are used for testing traffic. - Management—Network interfaces used by the Test Agents to connect to Routing Director. - Unknown—Network interfaces that the Test Agent is unable to identify or that have been removed. Based on your requirement, you can select or deselect Normal, Management, or Unknown. You can enable the Display Device Names toggle button to display the hostnames of the routers that are associated with the Test Agents. You can also enable the Hide unsupported toggle button to hide devices that are not supported by the plugin you have selected. To search for a specific interface or a Test Agent, enter one or more keywords in the search text box. You can search based on the name, description, IP and MAC address of the interfaces or the Test Agents. You can also search based on the device name, device model, and device MAC addresses if searching for a Test Agent associated with the device. You can also hover over an interface name to view the details such as the Test Agent name, bind

Table 298: TWAMP/TWAMP Light Parameters *(Continued)*

Parameter	Description
	interface, bind family, and to check if the Test Agent is run on a device or not. You can also remove the interface filters you have selected and reset the filters to the default mode by using the Reset filters button. The TWAMP/TWAMP Light plugin cannot be run on Test Agents associated with devices. To run TWAMP/TWAMP Light plugin on devices, use the RPM TWAMP/TWAMP Light plugin.

Table 298: TWAMP/TWAMP Light Parameters (Continued)

Parameter	Description
Reflectors*	Specify the reflector address. A Test Agent application can run a Reflector plugin whereas a Test Agent that is associated with a device needs to be configured to run Reflector plugin. When you click the text box, the Select reflectors page appears where you can add reflectors. On this page: • Reflector hostname—Specify the hostname for the reflector. The reflector receives the TWAMP request packets and reflect the packet back to the sender. Maximum length—64 characters • Test session port—Specify the destination port for the Test session. Range—0 through 65535 • Control session port—Specify the port for the control session. If set to zero, TWAMP Light is used. Default value—0 Range—0 through 65535 • Source port for Test session port—Specify the port for the control session. If set to zero, TWAMP Light is used. Specify the source port value for the Test session. Default value—0 Range—0 through 65535 • Name—This text box is automatically populated based on the data you specified in the Reflector hostname text box.

Table 298: TWAMP/TWAMP Light Parameters (*Continued*)

Parameter	Description
Rate	Specify the rate at which the TWAMP sender will send test traffic. The rate is calculated as the size of the request packet sent divided by the total request time. Each Ethernet packet contains one frame. Unit—Megabits per seconds (Mbit/s) Range—0.0 Mbit/s through 10000.0 Mbit/s
Time sync	Select the time synchronization option. Select Yes, if the clocks of Test Agent and the reflector are synchronized via Network Time Protocol (NTP) using external NTP servers. Select No to disable the time synchronization option. Select Inband to use the in-band synchronization algorithm.
Use hardware timestamping	Enable hardware timestamping, SO_TIMESTAMPING parameter, on TWAMP socket), if you want to use the network interface card (NIC) of Test Agents for delay and jitter measurements. If Test Agent NIC does not support it, an error message is displayed, and the measurement does not produce any delay metrics. By default, the toggle button is disabled.
Thresholds for error seconds (ES)	

Table 298: TWAMP/TWAMP Light Parameters (*Continued*)

Parameter	Description
Loss%	Specify the loss percentage value. If the loss percentage exceeds the configured value, an errored-second is raised. The Loss percentage is calculated by comparing the total number of packets that were lost with the total number of packets that were sent from the Test Agent. Unit—Percentage (%) Default value—0.0 % Range—0.0 % through 100.00 %
Delay	Specify the maximum tolerated value for delay. If there is time synchronization, the threshold value applies only to one-way delay. If there is no time synchronization, the threshold applies to round-trip delay. If delay exceeds the defined value, an errored-second is raised. Unit—Milliseconds (ms) Range—0.001 ms through 1000.0 ms
Delay Variation	Specify the maximum tolerated value for delay variance. If there is time synchronization, it applies to one-way delay. If there is no time synchronization, it applies to round-trip delay. If delay exceeds the defined value, an errored-second is raised. Unit—Milliseconds (ms) Range—0.001 ms through 1000.0 ms

Table 298: TWAMP/TWAMP Light Parameters (*Continued*)

Parameter	Description
Expected DSCP Value	Specify the expected DSCP or IP Precedence value on the IP packets received from the reflector device. If the received DSCP value does not match the configured value, an errored-second will be raised. Range—0 through 63
Thresholds for severely error seconds (SES)	
Loss	Specify the loss percentage value. If the loss percentage exceeds the configured value during a one-second interval, a severely errored-second is raised. The Loss percentage is calculated by comparing the total number of packets that were lost with the total number of packets that were sent from the Test Agent. Unit—Percentage (%) Minimum value—0.0 %
Delay	Specify the maximum acceptable delay. If there is time synchronization, it applies to one-way delay. If there is no time synchronization, it applies to round-trip delay. If delay exceeds the defined value, a severely errored-second is raised. Unit—Milliseconds (ms) Minimum value—0.001 ms

Table 298: TWAMP/TWAMP Light Parameters (*Continued*)

Parameter	Description
Delay variation	Specify the maximum tolerated value for delay variance (jitter). If there is time synchronization, it applies to one-way delay. If there is no time synchronization, it applies to round-trip delay. If delay exceeds the defined value, a severely errored-second is raised. Unit—Milliseconds (ms) Minimum value—0.001 ms
Advanced	
Frame Size	Specify the size of Layer 2 Ethernet frame size for the TWAMP test traffic. Frame Size indicates the total size of the data frame sent from the Test Agent to the remote endpoint. The size also includes the header size. Unit—Bytes Default value—1518 Range—87 through 9018
DSCP	Specify the Differentiated Services Code Point (DSCP) or IP Precedence value that is used in the IP packet headers. Default value—0 Range—0 through 63
Use random padding	Enable this to use random numbers or zeroes as padding in a TWAMP packet. By default, the toggle button is enabled.

Table 298: TWAMP/TWAMP Light Parameters *(Continued)*

Parameter	Description
Socket priority	Specify the value of SO_PRIORITY parameter of the TWAMP socket. Default value—0 Range—0 through 7
Socket send buffer size	Specify the value of SO_PRIORITY parameter of the TWAMP socket. Unit—Bytes Range—2048 through 10000000 bytes
Socket receive buffer size	Specify the value of SO_PRIORITY parameter of the TWAMP socket. Unit—Bytes Range—2048 through 10000000 bytes
Don't fragment flag	Enable the Don't Fragment Flag (DF Flag) to restrict the fragmentation of the packets that exceed the MTU. DF Flag is configured in an IP header. Router drops the packet if fragmentation is needed. Enabling this option may cause performance degradation both in the network and in the sending or receiving Test Agents. By default, the toggle button is enabled.

Table 298: TWAMP/TWAMP Light Parameters (*Continued*)

Parameter	Description
UAS period length	Specify the minimum value for the consecutive severely errored-seconds (SES) that causes a period of unavailability. The Unavailable Seconds (UAS) metric determines the number of seconds at which the service can be considered to be unavailable. Unit—Seconds (s) Default value—10 s Range—0 s through 300 s
Accept UDP checksum zero for IPv6	Enable Accept UDP Checksum Zero for IPv6 to accept the UDP Checksum as Zero for IPv6 in TWAMP Reflector packets. By default, the toggle button is enabled.
Percentiles	
First delay percentile	Specify the first delay percentile of the value of the TWAMP request packet. If the delay exceeds the configured value, the packet is included in the defined first delay percentile slot. Unit—Percentage (%) Range—0 % through 1 %
Second delay percentile	Specify the second delay percentile of the value of the TWAMP request packet. If the delay exceeds the configured value, the packet is included in the defined second delay percentile slot. Unit—Percentage (%) Range—0 % through 1 %

Table 298: TWAMP/TWAMP Light Parameters (*Continued*)

Parameter	Description
Threshold for first delay percentile	Specify the threshold for triggering an errored second based on the first delay percentile. Unit—Milliseconds (ms). Range—0.001 ms through 1000.00 ms
Threshold for second delay percentile	Specify the threshold for triggering an errored second based on the second delay percentile. Unit—Milliseconds (ms). Range—0.001 ms through 1000.00 ms
SES threshold for first delay percentile	Specify the threshold for triggering a severely errored second based on the first delay percentile. Unit—Milliseconds (ms). Range—0.001 ms through 1000.00 ms
SES threshold for second delay percentile	Specify the threshold for triggering a severely errored second based on the second delay percentile. Unit—Milliseconds (ms). Range—0.001 ms through 1000.00 ms
Periodic Streams— Periodic stream configuration allows you to generate the TWAMP traffic in bursts instead of generating the traffic continuously. This is generated in specific time period and the traffic generation can happen anytime in this time period.	

Table 298: TWAMP/TWAMP Light Parameters *(Continued)*

Parameter	Description
Active period duration	Specify the time duration of each cycle during which frames are sent. Ethernet Active period is followed by a silent period during which no ethernet frames are sent. Unit—Milliseconds (ms) Range—1 ms to 3600000 ms
Active cycle	Specify the time duration of the cycle starting with an active period and ending with a silent period. The Active cycle duration must be at least equal to the active period duration. Unit—Milliseconds (ms) Range—1 ms to 604800 ms
Report metrics during inactive period	Enable the toggle button to report metrics related to inactive periods of a periodic Test. By default, the toggle button is enabled.

Table 299: TWAMP Reflector Parameters

Parameter	Description
General	

Table 299: TWAMP Reflector Parameters *(Continued)*

Parameter	Description
Clients*	Select one or more Test Agent interfaces on which you want to run the Test. To select one or more Test Agent interfaces: 1. Click the Clients text box. The Test Agent/Interfaces choices page appears. 2. Select one or more Test Agent interfaces. For information on filtering Test Agent interfaces, see "Search for Test Agent interfaces by using Filters" on page 1195. 3. Click OK. NOTE: On the Test Agent/Interfaces choices page, you can filter the Test Agent interfaces based on the following criteria: • Interface name or the Test Agent name, or both. Based on your requirement, you can select or deselect Test Agents or Interfaces. • IPv4 address (with or without IPv4 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv4 or No IPv4 to view or hide the interfaces that are associated with the IPv4 family and the interfaces that are not associated with the IPv4 family. • IPv6 address (with or without IPv6 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv6 or No IPv6 to view or hide the interfaces that are associated with the IPv6 family and the interfaces that are not associated with the IPv6 family.

Table 299: TWAMP Reflector Parameters (Continued)

Parameter	Description
	- Deployment type of a Test Agent: - Application—Filter out the Test Agents deployed in a device as an application. - Device—Filter out the Test Agents associated with a device. Based on your requirement, you can select or deselect Application or Device to view or hide the Test Agents deployed in a device as an application or the Test Agents associated with a device. - Connection types between the interface and the router: - Link—Indicates that the status is Up and there is an active connection for the interface. - No Link—Indicates that the status is Down and there is no active connection for the interface. Based on your requirement, you can select or deselect Link or No Link. - Connection statuses of the Test Agent: - Online—The Test Agent is connected to Routing Director. - Offline—The Test Agent is not connected to Routing Director. When the interface supports RPM plugins, the interface contains <i>(On device)</i> at the end of its name indicating the Device Online status.

Table 299: TWAMP Reflector Parameters (Continued)

Parameter	Description
	When the interface supports only the normal plugins, the interface will not display (<i>On device</i>) at the end of its name, and the status refers to the Online status. Based on your requirement, you can select or deselect Online or Offline. - Interface types of the Test Agent: - Normal—Network interfaces that are used for testing traffic. - Management—Network interfaces used by the Test Agents to connect to Routing Director. - Unknown—Network interfaces that the Test Agent is unable to identify or that have been removed. Based on your requirement, you can select or deselect Normal, Management, or Unknown. You can enable the Display Device Names toggle button to display the hostnames of the routers that are associated with the Test Agents. You can also enable the Hide unsupported toggle button to hide devices that are not supported by the plugin you have selected. To search for a specific interface or a Test Agent, enter one or more keywords in the search text box. You can search based on the name, description, IP and MAC address of the interfaces or the Test Agents. You can also search based on the device name, device model, and device MAC addresses if searching for a Test Agent associated with the device. You can also hover over an interface name to view the details such as the Test Agent name, bind

Table 299: TWAMP Reflector Parameters *(Continued)*

Parameter	Description
	interface, bind family, and to check if the Test Agent is run on a device or not. You can also remove the interface filters you have selected and reset the filters to the default mode by using the Reset filters button.
Test session port	Specify the destination port on which the TWAMP reflector listens. Default value—7000 Range—1 through 65535
Rate threshold for ES	Specify the rate threshold for the combined rate for all sessions. If rate is below the threshold value, an errored-second is raised. Unit—Megabits per seconds (Mbit/s) Range—0.001 (Mbit/s) through 10000 (Mbit/s)
Standalone mode	Enable Standalone mode to push the metrics data to Routing Director. By default, the toggle button is disabled.

RPM Plug-ins

Table 300: RPM HTTP Parameters

Parameter	Description
General	

Table 300: RPM HTTP Parameters *(Continued)*

Parameter	Description
Clients*	Select one or more interfaces of the network devices on which you want to run a Test. To select one or more network devices as client values: 1. Click the Clients text box. The Test Agent/Interfaces choices page appears. 2. Select one or more Test Agent interfaces. For information on filtering Test Agent interfaces, see "Search for Test Agent interfaces by using Filters" on page 1199. 3. Click OK. NOTE: On the Test Agent/Interfaces choices page, you can filter the Test Agent interfaces based on the following criteria: • Interface name or the Test Agent name, or both. Based on your requirement, you can select or deselect Test Agents or Interfaces. NOTE: The Test Agent filter is not functional as the RPM plugins directly interact with the devices. • IPv4 address (with or without IPv4 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv4 or No IPv4 to view or hide the interfaces that are associated with the IPv4 family and the interfaces that are not associated with the IPv4 family. • IPv6 address (with or without IPv6 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv6 or No IPv6 to view or hide the

Table 300: RPM HTTP Parameters *(Continued)*

Parameter	Description
	interfaces that are associated with the IPv6 family and the interfaces that are not associated with the IPv6 family. • Deployment type of a Test Agent: • Application—Filter out the Test Agents deployed in a device as an application. • Device—Filter out the Test Agents associated with a device. Based on your requirement, you can select or deselect Application or Device to view or hide the Test Agents deployed in a device as an application or the Test Agents associated with a device. • Connection types between the interface and the router: • Link—Indicates that the status is Up and there is an active connection for the interface. • No Link—Indicates that the status is Down and there is no active connection for the interface. Based on your requirement, you can select or deselect Link or No Link. • Connection statuses of the Test Agent: • Online—The Test Agent is connected to Routing Director. • Offline—The Test Agent is not connected to Routing Director.

Table 300: RPM HTTP Parameters *(Continued)*

Parameter	Description
	When the interface supports RPM plugins, the interface contains <i>(On device)</i> at the end of its name indicating the Device Online status. When the interface supports only the normal plugins, the interface will not display <i>(On device)</i> at the end of its name, and the status refers to the Online status. Based on your requirement, you can select or deselect Online or Offline. - Interface types of the Test Agent: - Normal—Network interfaces that are used for testing traffic. - Management—Network interfaces used by the Test Agents to connect to Routing Director. - Unknown—Network interfaces that the Test Agent is unable to identify or that have been removed. Based on your requirement, you can select or deselect Normal, Management, or Unknown. You can enable the Display Device Names toggle button to display the hostnames of the routers that are associated with the Test Agents. You can also enable the Hide unsupported toggle button to hide devices that are not supported by the plugin you have selected. To search for a specific interface or a Test Agent, enter one or more keywords in the search text box. You can search based on the name, description, IP and MAC address of the interfaces or the Test Agents. You can also search based on the device name, device model, and device MAC addresses if

Table 300: RPM HTTP Parameters *(Continued)*

Parameter	Description
	searching for a Test Agent associated with the device. You can also hover over an interface name to view the details such as the Test Agent name, bind interface, bind family, and to check if the Test Agent is run on a device or not. You can also remove the interface filters you have selected and reset the filters to the default mode by using the Reset filters button.
URLs*	Specify the URL or host towards which the router sends HTTP Get requests. Maximum length—255 characters
Time between requests*	Specify the time taken between successive HTTP requests initiated by the router. Unit—Seconds (s) Default value—10 s Range—1 s to 255 s
Advanced	
Collection Interval*	Specify the collection interval. Collection interval is the frequency at which the results are collected from the device. Collection interval must be larger than the specified Time between requests. Unit—Seconds (s) Default value—15 s Range—5 s through 300 s

Table 300: RPM HTTP Parameters *(Continued)*

Parameter	Description
Device response timeout	Specify the device response timeout value. Device response timeout determines the time to wait for a response from the router. If the specified time exceeds, the connection to the device is considered to be failed. Device response timeout must be larger than the specified Collection Interval. Unit—Seconds (s) Default value—200 s Range—30 s through 300 s
Routing instance	Specify the name of routing instance on the router used by the probes.
TTL	Specify the number of times the packets hop before a device discards the packet. Time to live (TTL) value indicates lifespan of a request packet. TTL prevents infinite loop in a network when the packet fails to reach the destination. Default value—64. Range—1 through 254.
Metadata get	Enable Metadata get for the Test to perform HTTP Get requests for metadata to a target URL. By default, the toggle button is disabled.
IPv6 local link	Specify the router link-local logical interface name for the egress interface to be used. This applies only for Test sessions with IPv6 address as the target address. Maximum length—64 characters

Table 300: RPM HTTP Parameters *(Continued)*

Parameter	Description
Hardware timestamp	Specify the router link-local logical interface name for the egress interface to be used. This applies only for Test sessions with IPv6 address as the target address. By default, the toggle button is disabled.

Table 301: RPM PING Parameters

Parameter	Description
General	

Table 301: RPM PING Parameters (Continued)

Parameter	Description
Clients*	Select one or more interfaces of the network devices on which you want to run a Test. To select one or more network devices as client values: 1. Click the Clients text box. The Test Agent/Interfaces choices page appears. 2. Select one or more Test Agent interfaces. For information on filtering Test Agent interfaces, see "Search for Test Agent interfaces by using Filters" on page 1205. 3. Click OK. NOTE: On the Test Agent/Interfaces choices page, you can filter the Test Agent interfaces based on the following criteria: • Interface name or the Test Agent name, or both. Based on your requirement, you can select or deselect Test Agents or Interfaces. NOTE: The Test Agent filter is not functional as the RPM plugins directly interact with the devices. • IPv4 address (with or without IPv4 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv4 or No IPv4 to view or hide the interfaces that are associated with the IPv4 family and the interfaces that are not associated with the IPv4 family. • IPv6 address (with or without IPv6 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv6 or No IPv6 to view or hide the

Table 301: RPM PING Parameters *(Continued)*

Parameter	Description
	interfaces that are associated with the IPv6 family and the interfaces that are not associated with the IPv6 family. - Deployment type of a Test Agent: - Application—Filter out the Test Agents deployed in a device as an application. - Device—Filter out the Test Agents associated with a device. Based on your requirement, you can select or deselect Application or Device to view or hide the Test Agents deployed in a device as an application or the Test Agents associated with a device. - Connection types between the interface and the router: - Link—Indicates that the status is Up and there is an active connection for the interface. - No Link—Indicates that the status is Down and there is no active connection for the interface. Based on your requirement, you can select or deselect Link or No Link. - Connection statuses of the Test Agent: - Online—The Test Agent is connected to Routing Director. - Offline—The Test Agent is not connected to Routing Director.

Table 301: RPM PING Parameters *(Continued)*

Parameter	Description
	When the interface supports RPM plugins, the interface contains <i>(On device)</i> at the end of its name indicating the Device Online status. When the interface supports only the normal plugins, the interface will not display <i>(On device)</i> at the end of its name, and the status refers to the Online status. Based on your requirement, you can select or deselect Online or Offline. • Interface types of the Test Agent: • Normal—Network interfaces that are used for testing traffic. • Management—Network interfaces used by the Test Agents to connect to Routing Director. • Unknown—Network interfaces that the Test Agent is unable to identify or that have been removed. Based on your requirement, you can select or deselect Normal, Management, or Unknown. You can enable the Display Device Names toggle button to display the hostnames of the routers that are associated with the Test Agents. You can also enable the Hide unsupported toggle button to hide devices that are not supported by the plugin you have selected. To search for a specific interface or a Test Agent, enter one or more keywords in the search text box. You can search based on the name, description, IP and MAC address of the interfaces or the Test Agents. You can also search based on the device name, device model, and device MAC addresses if

Table 301: RPM PING Parameters *(Continued)*

Parameter	Description
	searching for a Test Agent associated with the device. You can also hover over an interface name to view the details such as the Test Agent name, bind interface, bind family, and to check if the Test Agent is run on a device or not. You can also remove the interface filters you have selected and reset the filters to the default mode by using the Reset filters button.
Hosts*	Specify the target hostname or IP address. When you click the text box, the Select Host page appears where you can enter hostnames. To add more than one Hosts, click + Add Host and specify the following: • Host—The hostname or the IP address of the remote endpoint. Maximum length—255 characters. • Name—The text box is automatically populated based on the data you specified in the Host text box.
Time between requests	Specify the time taken between successive ping requests initiated by the router. Unit—Seconds (s). Default value—10 s. Range—1 s through 255 s.
Advanced	

Table 301: RPM PING Parameters *(Continued)*

Parameter	Description
Collection interval	Specify the collection interval. Collection interval is the frequency at which the results are collected from the device. Collection interval must be larger than the specified Time between requests. Unit—Seconds (s). Default value—15 s. Range—5 s through 300 s.
Device response timeout	Device response timeout determines the time to wait for a response from the router. If the specified time is exceeded, the connection to the device is considered failed. Device response timeout must be larger than the specified Collection Interval. Unit—Seconds (s) Default value—200 s Range—30 s through 300 s
Routing instance	Specify the name of routing instance on the router used by the probes.
Data size	Specify the size of the data portion of the request packet sent from the router. Default value—0 Range—0 through 65400
Data fill	Specify the contents of the data portion of the request packets sent from the router. The value should be in hexadecimal format.

Table 301: RPM PING Parameters *(Continued)*

Parameter	Description
TTL	Specify the number of times the packets hop before a device discards the packet. Time to live (TTL) value indicates lifespan of a request packet. TTL prevents infinite loop in a network when the packet fails to reach the destination. For example, if you have entered the TTL value as 64, every time the packet passes a device, the value is reduced by one until the packet reaches the remote server. If the entered value reaches zero before reaching the remote endpoint, the packet is discarded. Default value—64 Range—1 through 254
DSCP code points	Specify the Differentiated Services Code Point (DSCP) value that is used in the IP packet headers. Maximum length—64 characters
Hardware timestamp	Enable hardware timestamping of RPM probe messages in the Packet Forwarding Engine host processor on the router. In case of ping plugin, you can enable timestamping of RPM messages in the Packet Forwarding Engine host processor. By default, the toggle button is disabled.
Ping timestamp	Enable ping timestamp to perform ICMP timestamp requests instead of ICMP echo requests. By default, the toggle button is disabled.
One way hardware timestamp	Enable the one-way hardware timestamping for one-way measurements (delay and jitter). By default, the toggle button is disabled.

Table 301: RPM PING Parameters *(Continued)*

Parameter	Description
IPv6 local link	Specify the router link-local logical interface name for the egress interface to be used. This applies only for Test sessions with IPv6 address as the target address. Maximum length—64 characters

Table 302: RPM TCP Parameters

Parameter	Description
General	

Table 302: RPM TCP Parameters (Continued)

Parameter	Description
Clients*	Select one or more interfaces of the network devices on which you want to run a Test. To select one or more network devices as client values: 1. Click the Clients text box. The Test Agent/Interfaces choices page appears. 2. Select one or more Test Agent interfaces. For information on filtering Test Agent interfaces, see "Search for Test Agent interfaces by using Filters" on page 1212. 3. Click OK. NOTE: On the Test Agent/Interfaces choices page, you can filter the Test Agent interfaces based on the following criteria: • Interface name or the Test Agent name, or both. Based on your requirement, you can select or deselect Test Agents or Interfaces. NOTE: The Test Agent filter is not functional as the RPM plugins directly interact with the devices. • IPv4 address (with or without IPv4 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv4 or No IPv4 to view or hide the interfaces that are associated with the IPv4 family and the interfaces that are not associated with the IPv4 family. • IPv6 address (with or without IPv6 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv6 or No IPv6 to view or hide the

Table 302: RPM TCP Parameters *(Continued)*

Parameter	Description
	interfaces that are associated with the IPv6 family and the interfaces that are not associated with the IPv6 family. - Deployment type of a Test Agent: - Application—Filter out the Test Agents deployed in a device as an application. - Device—Filter out the Test Agents associated with a device. Based on your requirement, you can select or deselect Application or Device to view or hide the Test Agents deployed in a device as an application or the Test Agents associated with a device. - Connection types between the interface and the router: - Link—Indicates that the status is Up and there is an active connection for the interface. - No Link—Indicates that the status is Down and there is no active connection for the interface. Based on your requirement, you can select or deselect Link or No Link. - Connection statuses of the Test Agent: - Online—The Test Agent is connected to Routing Director. - Offline—The Test Agent is not connected to Routing Director.

Table 302: RPM TCP Parameters *(Continued)*

Parameter	Description
	When the interface supports RPM plugins, the interface contains <i>(On device)</i> at the end of its name indicating the Device Online status. When the interface supports only the normal plugins, the interface will not display <i>(On device)</i> at the end of its name, and the status refers to the Online status. Based on your requirement, you can select or deselect Online or Offline. - Interface types of the Test Agent: - Normal—Network interfaces that are used for testing traffic. - Management—Network interfaces used by the Test Agents to connect to Routing Director. - Unknown—Network interfaces that the Test Agent is unable to identify or that have been removed. Based on your requirement, you can select or deselect Normal, Management, or Unknown. You can enable the Display Device Names toggle button to display the hostnames of the routers that are associated with the Test Agents. You can also enable the Hide unsupported toggle button to hide devices that are not supported by the plugin you have selected. To search for a specific interface or a Test Agent, enter one or more keywords in the search text box. You can search based on the name, description, IP and MAC address of the interfaces or the Test Agents. You can also search based on the device name, device model, and device MAC addresses if

Table 302: RPM TCP Parameters *(Continued)*

Parameter	Description
	searching for a Test Agent associated with the device. You can also hover over an interface name to view the details such as the Test Agent name, bind interface, bind family, and to check if the Test Agent is run on a device or not. You can also remove the interface filters you have selected and reset the filters to the default mode by using the Reset filters button.
Servers*	Specify the hostname or the IP address of the remote endpoint. When you click the text box, the Select server page appears where you can enter details. To add more than one servers, click + Add Server and specify the following: - Remote IP or hostname—The hostname or the IP address of the remote endpoint. Maximum length—255 characters - Remote TCP port—Specify the port number of TCP. Default value—7 Range—7 through 65535
Time between requests	Specify the time taken between successive TCP requests initiated by the router. Unit—Seconds (s) Default value—10 s Range—1 s through 255 s
Advanced	

Table 302: RPM TCP Parameters *(Continued)*

Parameter	Description
Collection interval	Specify the collection interval. Collection interval is the frequency at which the results are collected from the device. Collection interval must be larger than the specified Time between requests. Unit—Seconds (s) Default value—15 s Range—5 s through 300 s
Device response timeout	Specify the device response timeout value. Device response timeout determines the time to wait for a response from the router. If the specified time is exceeded, the connection to the device is considered failed. Device response timeout must be larger than the specified Collection Interval. Unit—Seconds (s) Default value—200 s Range—30 s through 300 s
Routing instance	Specify the name of routing instance on the router used by the request packets. Maximum value—64
Data size	Specify the size of the data portion of the request packets sent from the router. Default value—0 Range—0 through 65400
Data fill	Specify the contents of the data portion of the request packets sent from the router. The value should be in hexadecimal format.

Table 302: RPM TCP Parameters *(Continued)*

Parameter	Description
TTL	Specify the number of times the packets hop before a device discards the packet. Time to live (TTL) value indicates lifespan of a request packet. TTL prevents infinite loop in a network when the packet fails to reach the destination. For example, if you have entered the TTL value as 64, every time the packet passes a device, the value is reduced by one until the packet reaches the remote server. If the entered value reaches zero before reaching the remote endpoint, the packet is discarded. Default value—64 Range—1 through 254
DSCP code points	Specify the Differentiated Services Code Point (DSCP) value that is used in the IP packet headers. Maximum length—64 characters
IPv6 Local Link	Specify the router link-local logical interface name for the egress interface to be used. This applies only for Test sessions with IPv6 address as the target address. Maximum length—64 characters

Table 303: RPM TWAMP Parameters

Parameter	Description
General	

Table 303: RPM TWAMP Parameters *(Continued)*

Parameter	Description
Clients*	Select one or more interfaces of the network devices on which you want to run a Test. To select one or more network devices as client values: 1. Click the Clients text box. The Test Agent/Interfaces choices page appears. 2. Select one or more Test Agent interfaces. For information on filtering Test Agent interfaces, see "Search for Test Agent interfaces by using Filters" on page 1218. 3. Click OK. NOTE: On the Test Agent/Interfaces choices page, you can filter the Test Agent interfaces based on the following criteria: • Interface name or the Test Agent name, or both. Based on your requirement, you can select or deselect Test Agents or Interfaces. NOTE: The Test Agent filter is not functional as the RPM plugins directly interact with the devices. • IPv4 address (with or without IPv4 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv4 or No IPv4 to view or hide the interfaces that are associated with the IPv4 family and the interfaces that are not associated with the IPv4 family. • IPv6 address (with or without IPv6 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv6 or No IPv6 to view or hide the

Table 303: RPM TWAMP Parameters *(Continued)*

Parameter	Description
	interfaces that are associated with the IPv6 family and the interfaces that are not associated with the IPv6 family. - Deployment type of a Test Agent: - Application—Filter out the Test Agents deployed in a device as an application. - Device—Filter out the Test Agents associated with a device. Based on your requirement, you can select or deselect Application or Device to view or hide the Test Agents deployed in a device as an application or the Test Agents associated with a device. - Connection types between the interface and the router: - Link—Indicates that the status is Up and there is an active connection for the interface. - No Link—Indicates that the status is Down and there is no active connection for the interface. Based on your requirement, you can select or deselect Link or No Link. - Connection statuses of the Test Agent: - Online—The Test Agent is connected to Routing Director. - Offline—The Test Agent is not connected to Routing Director.

Table 303: RPM TWAMP Parameters *(Continued)*

Parameter	Description
	When the interface supports RPM plugins, the interface contains <i>(On device)</i> at the end of its name indicating the Device Online status. When the interface supports only the normal plugins, the interface will not display <i>(On device)</i> at the end of its name, and the status refers to the Online status. Based on your requirement, you can select or deselect Online or Offline. • Interface types of the Test Agent: • Normal—Network interfaces that are used for testing traffic. • Management—Network interfaces used by the Test Agents to connect to Routing Director. • Unknown—Network interfaces that the Test Agent is unable to identify or that have been removed. Based on your requirement, you can select or deselect Normal, Management, or Unknown. You can enable the Display Device Names toggle button to display the hostnames of the routers that are associated with the Test Agents. You can also enable the Hide unsupported toggle button to hide devices that are not supported by the plugin you have selected. To search for a specific interface or a Test Agent, enter one or more keywords in the search text box. You can search based on the name, description, IP and MAC address of the interfaces or the Test Agents. You can also search based on the device name, device model, and device MAC addresses if

Table 303: RPM TWAMP Parameters *(Continued)*

Parameter	Description
	searching for a Test Agent associated with the device. You can also hover over an interface name to view the details such as the Test Agent name, bind interface, bind family, and to check if the Test Agent is run on a device or not. You can also remove the interface filters you have selected and reset the filters to the default mode by using the Reset filters button.

Table 303: RPM TWAMP Parameters *(Continued)*

Parameter	Description
Hosts*	Specify IPv4 or IPv6 Address of TWAMP reflector. When you click the text box, the Select reflectors page appears where you can add reflectors. On this page: • Reflector hostname—Specify the hostname for the reflector. Maximum length—64 characters • Test session port—Specify IPv4 or IPv6 Address of TWAMP reflector. Range—0 through 65535 • Control session port—Specify the port for the control session. If the port value is zero, it is for TWAMP Light, and if port value is in the range of 862-65535, it is for managed TWAMP. Default value—0 Range—0 through 65535 • Source port for Test session port—Specify the source port value for the Test session. Default value—0 Range—0 through 65535 • Name—The text box is automatically populated based on the data you specified in the Reflector hostname text box.
Time between requests	Specify the time taken between successive TWAMP requests initiated by the router. Unit—Seconds (s) Default value—10 s Range—1 s through 255 s

Table 303: RPM TWAMP Parameters *(Continued)*

Parameter	Description
Advanced	
Collection interval	Specify the collection interval. Collection interval is the frequency at which the results are collected from the device. Collection interval must be larger than the specified Time between requests. Unit—Seconds (s) Default value—15 s Range—5 s through 300 s
Device response timeout	Specify the device response timeout value. Device response timeout determines the time to wait for a response from the router. If the specified time is exceeded, the connection to the device is considered failed. Device response timeout must be larger than the specified Collection Interval. Unit—Seconds (s) Default value—200 s Range—30 s through 300 s
Routing instance	Specify the name of routing instance on the router used by the request packets. Maximum value—64
Zero fill	Enable Zero Fill to populate the content for the request packet with zeros.

Table 303: RPM TWAMP Parameters *(Continued)*

Parameter	Description
Data size	Specify the size of the data portion of the request packets sent from the router. Default value—60. Range—60 to 1400.
TTL	Specify the number of times the packets hop before a device discards the packet. Time to live (TTL) value indicates lifespan of a request packet. TTL prevents infinite loop in a network when the packet fails to reach the destination. For example, if you have entered the TTL value as 64, every time the packet passes a device, the value is reduced by one until the packet reaches the remote server. If the entered value reaches zero before reaching the remote endpoint, the packet is discarded. Default value—64 Range—1 to 254
DSCP Code Points	Specify the Differentiated Services Code Point (DSCP) value that is used in the IP packet headers. Maximum length—64 characters
IPv6 Local Link	Specify the router link-local logical interface name for the egress interface to be used. This applies only for Test sessions with IPv6 address as the target address. Maximum length—64 characters
PFE timestamping	Enable PFE timestamping to perform timestamping on Packet Forward Engine host. By default, the toggle button is disabled.

Table 304: RPM UDP Parameters

Parameter	Description
General	

Table 304: RPM UDP Parameters *(Continued)*

Parameter	Description
Clients*	Select one or more interfaces of the network devices on which you want to run a Test. To select one or more network devices as client values: 1. Click the Clients text box. The Test Agent/Interfaces choices page appears. 2. Select one or more Test Agent interfaces. For information on filtering Test Agent interfaces, see "Search for Test Agent interfaces by using Filters" on page 1226. 3. Click OK. NOTE: On the Test Agent/Interfaces choices page, you can filter the Test Agent interfaces based on the following criteria: • Interface name or the Test Agent name, or both. Based on your requirement, you can select or deselect Test Agents or Interfaces. NOTE: The Test Agent filter is not functional as the RPM plugins directly interact with the devices. • IPv4 address (with or without IPv4 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv4 or No IPv4 to view or hide the interfaces that are associated with the IPv4 family and the interfaces that are not associated with the IPv4 family. • IPv6 address (with or without IPv6 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv6 or No IPv6 to view or hide the

Table 304: RPM UDP Parameters *(Continued)*

Parameter	Description
	interfaces that are associated with the IPv6 family and the interfaces that are not associated with the IPv6 family. • Deployment type of a Test Agent: • Application—Filter out the Test Agents deployed in a device as an application. • Device—Filter out the Test Agents associated with a device. Based on your requirement, you can select or deselect Application or Device to view or hide the Test Agents deployed in a device as an application or the Test Agents associated with a device. • Connection types between the interface and the router: • Link—Indicates that the status is Up and there is an active connection for the interface. • No Link—Indicates that the status is Down and there is no active connection for the interface. Based on your requirement, you can select or deselect Link or No Link. • Connection statuses of the Test Agent: • Online—The Test Agent is connected to Routing Director. • Offline—The Test Agent is not connected to Routing Director.

Table 304: RPM UDP Parameters *(Continued)*

Parameter	Description
	When the interface supports RPM plugins, the interface contains <i>(On device)</i> at the end of its name indicating the Device Online status. When the interface supports only the normal plugins, the interface will not display <i>(On device)</i> at the end of its name, and the status refers to the Online status. Based on your requirement, you can select or deselect Online or Offline. - Interface types of the Test Agent: - Normal—Network interfaces that are used for testing traffic. - Management—Network interfaces used by the Test Agents to connect to Routing Director. - Unknown—Network interfaces that the Test Agent is unable to identify or that have been removed. Based on your requirement, you can select or deselect Normal, Management, or Unknown. You can enable the Display Device Names toggle button to display the hostnames of the routers that are associated with the Test Agents. You can also enable the Hide unsupported toggle button to hide devices that are not supported by the plugin you have selected. To search for a specific interface or a Test Agent, enter one or more keywords in the search text box. You can search based on the name, description, IP and MAC address of the interfaces or the Test Agents. You can also search based on the device name, device model, and device MAC addresses if

Table 304: RPM UDP Parameters *(Continued)*

Parameter	Description
	searching for a Test Agent associated with the device. You can also hover over an interface name to view the details such as the Test Agent name, bind interface, bind family, and to check if the Test Agent is run on a device or not. You can also remove the interface filters you have selected and reset the filters to the default mode by using the Reset filters button.
Servers*	Specify the remote IP address or hostname of the server to which Test Agents send the requests. When you click the text box, the Select server page appears where you can enter details. To add more than one servers, click + Add Server and specify the following: • Servers—Specify the hostname or the IP address of the remote endpoint. Maximum length—255 characters. • Remote TCP port—Specify the port number of UDP. Default value—7. Range—7 through 65535.
Time Between Requests	Specify the time taken between successive UDP requests initiated by the routers. Unit—Seconds (s). Default value—10 s. Range—1 s through 255 s.

Table 304: RPM UDP Parameters (Continued)

Parameter	Description
Remote port	Configure the remote port number for the Test sessions. Default value—7. Range—7 through 65535.
Advanced	
Collection Interval	Specify the collection interval. Collection interval is the frequency at which the results are collected from the remote endpoint or the device. Collection interval must be larger than the specified Time between requests. Unit—Seconds (s). Default value—15 s. Range—5 s through 300 s.
Device Response Timeout	Specify the device response timeout value. Device response timeout determines the time to wait for a response from the router. If the specified time is exceeded, the connection to the device is considered failed. Device response timeout must be larger than the specified Collection Interval. Unit—Seconds (s). Default value—200 s. Range—30 s through 300 s.
Routing instance	Specify the name of routing instance on the router used by the probes. Maximum value—64.

Table 304: RPM UDP Parameters *(Continued)*

Parameter	Description
Data Size	Specify the size of the data portion of the request packets sent from the router. Default value—0. Range—0 through 65400.
Data fill	Specify the contents of the data portion of the request packets sent from the router. The value should be in hexadecimal format. The value should be in hexadecimal format.
TTL	Specify the number of times the packets hop before a device discards the packet. Time to live (TTL) value indicates lifespan of a request packet. TTL prevents infinite loop in a network when the packet fails to reach the destination. For example, if you have entered the TTL value as 64, every time the packet passes a device, the value is reduced by one until the packet reaches the remote server. If the entered value reaches zero before reaching the remote endpoint, the packet is discarded. Default value—64. Range—1 through 254.
DSCP Code Points	Specify the Differentiated Services Code Point (DSCP) value that is used in the IP packet headers. Maximum length—64 characters.
Hardware Timestamp	Enable hardware timestamping of RPM probe messages in the Packet Forwarding Engine host processor on the router. By default, the toggle button is disabled.

Table 304: RPM UDP Parameters *(Continued)*

Parameter	Description
Ping Timestamp	Enable ping timestamping to send UDP timestamp requests. By default, the toggle button is disabled.
One Way Hardware Timestamp	Enable the one-way hardware timestamping for one-way delay and jitter measurements. By default, the toggle button is disabled.
IPv6 Local Link	Specify the router link-local logical interface name for the egress interface to be used. This applies only for Test sessions with IPv6 address as the target address. Maximum length—64 characters.

Table 305: Y.1731 DM Parameters

Parameter	Description
General	

Table 305: Y.1731 DM Parameters (Continued)

Parameter	Description
Senders*	Select one or more interfaces of the network devices on which you want to run a Monitor. To select one or more network devices as client values: 1. Click the Clients text box. The Test Agent/Interfaces choices page appears. 2. Select one or more Test Agent interfaces. For information on filtering Test Agent interfaces, see "Search for Test Agent interfaces by using Filters" on page 1233. 3. Click OK. NOTE: On the Test Agent/Interfaces choices page, you can filter the Test Agent interfaces based on the following criteria: • Interface name or the Test Agent name, or both. Based on your requirement, you can select or deselect Test Agents or Interfaces. NOTE: The Test Agent filter is not functional as the RPM plugins directly interact with the devices. • IPv4 address (with or without IPv4 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv4 or No IPv4 to view or hide the interfaces that are associated with the IPv4 family and the interfaces that are not associated with the IPv4 family. • IPv6 address (with or without IPv6 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv6 or No IPv6 to view or hide the

Table 305: Y.1731 DM Parameters *(Continued)*

Parameter	Description
	interfaces that are associated with the IPv6 family and the interfaces that are not associated with the IPv6 family. - Deployment type of a Test Agent: - Application—Filter out the Test Agents deployed in a device as an application. - Device—Filter out the Test Agents associated with a device. Based on your requirement, you can select or deselect Application or Device to view or hide the Test Agents deployed in a device as an application or the Test Agents associated with a device. - Connection types between the interface and the router: - Link—Indicates that the status is Up and there is an active connection for the interface. - No Link—Indicates that the status is Down and there is no active connection for the interface. Based on your requirement, you can select or deselect Link or No Link. - Connection statuses of the Test Agent: - Online—The Test Agent is connected to Routing Director. - Offline—The Test Agent is not connected to Routing Director.

Table 305: Y.1731 DM Parameters *(Continued)*

Parameter	Description
	When the interface supports RPM plugins, the interface contains <i>(On device)</i> at the end of its name indicating the Device Online status. When the interface supports only the normal plugins, the interface will not display <i>(On device)</i> at the end of its name, and the status refers to the Online status. Based on your requirement, you can select or deselect Online or Offline. - Interface types of the Test Agent: - Normal—Network interfaces that are used for testing traffic. - Management—Network interfaces used by the Test Agents to connect to Routing Director. - Unknown—Network interfaces that the Test Agent is unable to identify or that have been removed. Based on your requirement, you can select or deselect Normal, Management, or Unknown. You can enable the Display Device Names toggle button to display the hostnames of the routers that are associated with the Test Agents. You can also enable the Hide unsupported toggle button to hide devices that are not supported by the plugin you have selected. To search for a specific interface or a Test Agent, enter one or more keywords in the search text box. You can search based on the name, description, IP and MAC address of the interfaces or the Test Agents. You can also search based on the device name, device model, and device MAC addresses if

Table 305: Y.1731 DM Parameters *(Continued)*

Parameter	Description
	searching for a Test Agent associated with the device. You can also hover over an interface name to view the details such as the Test Agent name, bind interface, bind family, and to check if the Test Agent is run on a device or not. You can also remove the interface filters you have selected and reset the filters to the default mode by using the Reset filters button. NOTE: The filters IPv4 and IPv6 are not applicable in the case of Y.1731 SLM.

Table 305: Y.1731 DM Parameters (*Continued*)

Parameter	Description
SLA Iterators*	Specify one or more performance monitoring iterator profile names. When you click the text box, the Select iterator page appears where you can enter details. To add more than one iterator, click + Add iterator and specify the following: • Iterator profile name—Specify the iterator profile name used for SLA performance monitoring. • Maintenance domain name—Specify the maintenance domain name. A maintenance domain is a unique identifier for managing a network within a specific boundary. This domain name represents the space that is owned and operated by a single entity, which is defined by the set of internal ports and boundary ports. • Maintenance association name—Specify the maintenance association name. A maintenance association identifies services within a maintenance domain, and it logically groups the endpoints involved in monitoring the same service. The CMF protocol operates within this association. • Local maintenance end point—Specify the local maintenance endpoint (MEP). A local MEP is a device or a port within a maintenance domain. It is responsible for initiating a measurement to the remote MEP as well as detecting issues within a maintenance domain. • Remote maintenance end point—Specify the local maintenance endpoint (MEP). A local MEP is a device or a port within a maintenance domain. It is responsible for initiating

Table 305: Y.1731 DM Parameters *(Continued)*

Parameter	Description
	a measurement to the remote MEP as well as detecting issues within a maintenance domain. • Name—Specify a unique identifying name for the SLA iterator.
Collection interval (s)	Specify the collection interval. Collection interval is the frequency at which the results related to delay and jitter are collected from the router. Unit—Seconds (s). Default value—15 s. Range—5 s through 300 s.
Device response timeout	Specify the device response timeout value. Device response timeout determines the time to wait for a response from the router. If the specified time exceeds then the connection to the router is considered to be failed. Device response timeout must be larger than the specified Collection Interval. Unit—Seconds (s). Default value—200 s. Range—30 s through 300 s.

Table 306: Y.1731 SLM Parameters

Parameter	Description
General	

Table 306: Y.1731 SLM Parameters *(Continued)*

Parameter	Description
Senders*	Select one or more interfaces of the network devices on which you want to run a Monitor. To select one or more network devices as client values: 1. Click the Clients text box. The Test Agent/Interfaces choices page appears. 2. Select one or more Test Agent interfaces. For information on filtering Test Agent interfaces, see "Search for Test Agent interfaces by using Filters" on page 1239. 3. Click OK. NOTE: On the Test Agent/Interfaces choices page, you can filter the Test Agent interfaces based on the following criteria: • Interface name or the Test Agent name, or both. Based on your requirement, you can select or deselect Test Agents or Interfaces. NOTE: The Test Agent filter is not functional as the RPM plugins directly interact with the devices. • IPv4 address (with or without IPv4 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv4 or No IPv4 to view or hide the interfaces that are associated with the IPv4 family and the interfaces that are not associated with the IPv4 family. • IPv6 address (with or without IPv6 addresses) of a Test Agent. Based on your requirement, you can select or deselect IPv6 or No IPv6 to view or hide the

Table 306: Y.1731 SLM Parameters *(Continued)*

Parameter	Description
	interfaces that are associated with the IPv6 family and the interfaces that are not associated with the IPv6 family. - Deployment type of a Test Agent: - Application—Filter out the Test Agents deployed in a device as an application. - Device—Filter out the Test Agents associated with a device. Based on your requirement, you can select or deselect Application or Device to view or hide the Test Agents deployed in a device as an application or the Test Agents associated with a device. - Connection types between the interface and the router: - Link—Indicates that the status is Up and there is an active connection for the interface. - No Link—Indicates that the status is Down and there is no active connection for the interface. Based on your requirement, you can select or deselect Link or No Link. - Connection statuses of the Test Agent: - Online—The Test Agent is connected to Routing Director. - Offline—The Test Agent is not connected to Routing Director.

Table 306: Y.1731 SLM Parameters *(Continued)*

Parameter	Description
	When the interface supports RPM plugins, the interface contains <i>(On device)</i> at the end of its name indicating the Device Online status. When the interface supports only the normal plugins, the interface will not display <i>(On device)</i> at the end of its name, and the status refers to the Online status. Based on your requirement, you can select or deselect Online or Offline. • Interface types of the Test Agent: • Normal—Network interfaces that are used for testing traffic. • Management—Network interfaces used by the Test Agents to connect to Routing Director. • Unknown—Network interfaces that the Test Agent is unable to identify or that have been removed. Based on your requirement, you can select or deselect Normal, Management, or Unknown. You can enable the Display Device Names toggle button to display the hostnames of the routers that are associated with the Test Agents. You can also enable the Hide unsupported toggle button to hide devices that are not supported by the plugin you have selected. To search for a specific interface or a Test Agent, enter one or more keywords in the search text box. You can search based on the name, description, IP and MAC address of the interfaces or the Test Agents. You can also search based on the device name, device model, and device MAC addresses if

Table 306: Y.1731 SLM Parameters *(Continued)*

Parameter	Description
	searching for a Test Agent associated with the device. You can also hover over an interface name to view the details such as the Test Agent name, bind interface, bind family, and to check if the Test Agent is run on a device or not. You can also remove the interface filters you have selected and reset the filters to the default mode by using the Reset filters button. NOTE: The filters IPv4 and IPv6 are not applicable in the case of Y.1731 SLM.

Table 306: Y.1731 SLM Parameters (Continued)

Parameter	Description
SLA Iterators*	Specify one or more performance monitoring iterator profile names. When you click the text box, the Select iterator page appears where you can enter details. To add more than one iterator, click + Add iterator and specify the following: • Iterator profile name—Specify the iterator profile name used for SLA performance monitoring. • Maintenance domain name—Specify the maintenance domain name. A maintenance domain is a unique identifier for managing a network within a specific boundary. This domain name represents the space that is owned and operated by a single entity, which is defined by the set of internal ports and boundary ports. • Maintenance association name—Specify the maintenance association name. A maintenance association identifies services within a maintenance domain, and it logically groups the endpoints involved in monitoring the same service. The CMF protocol operates within this association. • Local maintenance end point—Specify the local maintenance endpoint (MEP). A local MEP is a device or a port within a maintenance domain. It is responsible for initiating a measurement to the remote MEP as well as detecting issues within a maintenance domain. • Remote maintenance end point—Specify the local maintenance endpoint (MEP). A local MEP is a device or a port within a maintenance domain. It is responsible for initiating

Table 306: Y.1731 SLM Parameters *(Continued)*

Parameter	Description
	a measurement to the remote MEP as well as detecting issues within a maintenance domain. • Name—Specify a unique identifying name for the SLA iterator.
Collection interval (s)	Specify the collection interval. Collection interval is the frequency at which the results related to delay and jitter are collected from the router. Unit—Seconds (s). Default value—15 s. Range—5 s through 300 s.
Device response timeout	Specify the device response timeout value. Device response timeout determines the time to wait for a response from the router. If the specified time exceeds then the connection to the router is considered to be failed. Device response timeout must be larger than the specified Collection Interval. Unit—Seconds (s). Default value—200 s. Range—30 s through 300 s.

Configure Evaluation Criteria

Table 307: Evaluation Criteria

Field	Description
Field	Select the type of metric that you want to evaluate from the drop-down list. The metrics listed in this drop-down depends on the Task you select. The metrics will be displayed as Stream graphs in the <i>Stream-Name</i> Details page.
Comparator	Select the type of comparator that you want to use for the evaluation. You can choose among the following comparators— ==(equal to), != (not equal to), < (less than), <= (less than or equal to), > (greater than), >= (greater than or equal to). For example, if you have used > (greater than) comparator, and if you have specified 3000 ms in the Value text box, then an event is raised if the value exceeds 3000 ms.
Value	Specify the threshold value that you want to evaluate. On violation of any configured threshold value, an event is generated.

Table 307: Evaluation Criteria *(Continued)*

Field	Description
Severity	Select the severity level of the event that is raised when there is a threshold violation. The severity levels are: Debug—Indicates that the event needs troubleshooting. Info—Indicates that it is an informational event and does not require immediate attention. Warning—Indicates that the event needs to be reviewed but does not require immediate attention. Error—Indicates that the event needs to be reviewed and requires immediate attention. Critical—Indicates that the event is major and needs immediate attention. When an event is raised, the severity level you have selected will be displayed in the Events bar of the <i>Monitor-Name</i> page.
Name	Specify the name for the evaluation criteria.
Description	Specify the description to be displayed when there is a threshold violation. If you do not specify the description, the default description generated by Routing Director is displayed. When an event is raised, the description you have specified here will be displayed in the Events bar of the <i>Monitor-Name</i> page.
Raise Delay	Specify the duration that the Routing Director waits before triggering an event. If the number of seconds exceeds the configured Raise delay value, then an event is triggered. The triggered event is displayed in the Events bar of the <i>Monitor-Name</i> page.

Table 307: Evaluation Criteria *(Continued)*

Field	Description
Clear Delay	Specify the duration of time after which an event is cleared if the threshold is not violated. If the threshold has not exceeded the configured Clear delay seconds, then the event is cleared.

RELATED DOCUMENTATION

- [About the Monitors Page | 1247](#)
- [Active Assurance Workflow | 891](#)

About the Monitors Page

IN THIS SECTION

- [Tasks You Can Perform | 1248](#)

To access this page, click **Observability > Active Assurance > Monitors**.

Monitors are entities that you configure to indefinitely measure metrics (such as response time and packet loss). A Monitor contains only one Step and a Step can have multiple Tasks. For more information on Monitors, see "[Tests and Monitors Overview](#)" on page 944.

You (superusers and network administrators) can use the Monitors page to view the list of Monitors that the Test Agents run in your network and view the details of each Monitor. You can also view the following widgets on the Monitors page:

- Urgent actions needed—Displays the total number of Monitors that need immediate action or intervention to resolve one or more critical issues.
- Actions needed—Displays the total number of Monitors that need action or intervention to resolve major issues, but not immediately.

- **Total**—Displays the total number of Monitors that you have created in the network.
- **Running**—Displays the total number of Monitors that are running in the network.

Tasks You Can Perform

You can perform the following tasks on the Monitors page:

- **View details of a Monitor**—You can view the list of all the Monitors in your network and the details of these Monitors.

To view the details of a Monitor, select a Monitor and click **More > Details**. You can also hover over the *Monitor-Name* and click the **Detailed View** icon. The Monitor Details pane appears on the right side of the page displaying the Monitor details.

On the Monitor Details pane, you can:

- **Copy API Request URL**—You can copy the API Request URL so that you can reuse the copied API request URL to fetch the Monitor details when you restart the Monitor.
- View information such as, name, description, tags, ID, creation, and execution status of the Monitor.
- Edit the name, description, and tags of the Monitor.

Click the Close (x) icon to close the pane.

- **View details of a Monitor**—You can view the list of all the Monitors that you have run and the details of a selected Monitor. [Table 308 on page 1249](#) describes the fields on the Monitors page.

Click a *Monitor-Name* to view the details of the Monitor. For more information, see ["About the Monitor-Name Page" on page 1253](#).

- Create a Monitor. See ["Create a Monitor" on page 1110](#).
- You can also perform the following tasks on this page:
 - Sort, resize, or re-arrange columns in a table (grid).



NOTE: Sort functionality does not work for Created By, Updated By, Last Started By, Last Stopped By and Tags.

- Show or hide columns in the table or reset page preferences, using the vertical ellipsis menu.

- Filter the data displayed in the table—Click the filter icon (funnel) and select whether you want to show or hide advanced filters. You can then add or remove filter criteria, save criteria as a filter, apply or clear filters, and so on. The filtered results are displayed on the same page.

For more information, see ["GUI Overview" on page 6](#).

Table 308: Fields on the Monitors Page

Fields	Description
Name	The name of the Monitor that you specified when you created the Monitor.
Description	The description of the Monitor that you specified when you created the Monitor.
Status	The statuses for the Monitor are: • Started—The Test Agent has initiated a Monitor. • Running—The Test Agent is running the Monitor. • Error—The Test Agent has encountered an error while executing the Monitor. The errors can be anything from a Test Agent being offline to a metric configuration failure or a timeout. • Stopped—The Test Agent has stopped running the Monitor.
Requested Status	The expected status of the Monitor.
Created By	The username (e-mail address) of the user who has created the Monitor.
Updated By	The username (e-mail address) of the user who has updated the Monitor.
Created	The date and time when the Monitor was created. The timestamp is displayed in the following format: Month DD, YYYY, HH:MM:SS AM/PM. For example, Feb 5, 2024, 4:29:52 PM.

Table 308: Fields on the Monitors Page (*Continued*)

Fields	Description
Last Started	The date and time when the Monitor was last started. The timestamp is displayed in the following format: Month DD, YYYY, HH:MM:SS AM/PM. For example, Feb 5, 2024, 4:29:52 PM.
Last Updated	The latest date and time when the Monitor was last updated. The timestamp is displayed in the following format: Month DD, YYYY, HH:MM:SS AM/PM. For example, Feb 5, 2024, 4:29:52 PM.
Tags	The tags you have configured for the Monitor in the <i>key:value</i> format. By default, only two tags are displayed in column. To view more than two tags, click + Number more. A tag is a key-value pair in which the key signifies a category for which you configure a value. The value is an identifier for the category. Examples of key-value pairs are device and device name (edgedevice:acx7000), site and site name (site:Bengaluru). You can configure tags for a Monitor at the time of creating a Monitor.
Last Started By	The username (e-mail address) of the user who started the Monitor lately.
Last Stopped	The date and time when the Monitor was last stopped. The timestamp is displayed in the following format: Month DD, YYYY, HH:MM:SS AM/PM. For example, Feb 5, 2024, 4:29:52 PM.
Last Stopped By	The username (e-mail address) of the user who has stopped the Monitor lately.

Table 308: Fields on the Monitors Page (*Continued*)

Fields	Description
Events	The total number of Monitor-related events generated by Routing Director. The Events column display the number of issues that are related to Monitors or measurements related to Monitors. It also displays icons representing different severity levels in the order of priority. The highest priority is marked as Critical, indicating the most urgent issues, while the lowest priority is labeled as Informational, signifying minor notifications. Click the <i>Event-Number-link</i> and the Events page appears displaying the list of Monitor-related events generated by Routing Director. For more information, see Table 309 on page 1251.

Table 309: Description of fields on the Events Page

Field	Description
Severity	The severity level of the event that is raised by Test Agents. The following are the severity levels: Critical—Indicates that the event is critical and the issue needs immediate attention. Error—Indicates that the issue that has caused the event needs to be fixed and requires immediate attention and troubleshooting. Warning—Indicates that the issue that has caused the event needs to be fixed but does not require immediate attention. Info—Indicates the progress of a task and provides information on the event. The informational event does not require immediate attention.
Description	The description that you specified when you configured the evaluation criteria for a Test.

Table 309: Description of fields on the Events Page *(Continued)*

Field	Description
Raise Time	The date and time when the event was generated. The timestamp is displayed in the following format: Month DD, YYYY, HH:MM:SS AM/PM. The date and time is displayed according to the Raise delay you specified when you configured the evaluation criteria for a Test. For example, Sep 5, 2024, 4:29:52 PM.
Clear Time	The date and time when the event was cleared. The timestamp is displayed in the following format: Month DD, YYYY, HH:MM:SS AM/PM. The date and time is displayed according to the Clear delay you specified when you configured the evaluation criteria for a Test. For example, Sep 5, 2024, 4:29:52 PM.
Subject	The details of the event generated is displayed in JSON format. It displays various IDs that are related to the event like Test Agent ID, Test ID, Task ID, Stream ID, and so on. To view the details, click show hyperlink.
Data	The details of the evaluation criteria associated with the event generated. To view data, click show hyperlink.

RELATED DOCUMENTATION

| [About the Test-Name Page](#) | 1104

About the *Monitor-Name* Page

IN THIS SECTION

- [Tasks You Can Perform](#) | 1253

To access the *Monitor-Name* page:

1. Select **Observability > Active Assurance > Monitors**.

The Monitors page appears.

2. Click a *Monitor-Name*.

The *Monitor-Name* page appears displaying the Monitor results.

You (superusers and network administrators) can use this page to view the details of the Monitor that you have selected. You can view general details of a Monitor such as the Monitor name, Monitor description, time range of the Monitor, the overall status of the Monitor, and the list of events generated. You can also view the Streams generated for each Task that you have configured for the Monitor.

Tasks You Can Perform

You can perform the following tasks on the *Monitor-Name* page:

- View Monitor details—You can view the following details related to the Monitor:
 - Monitor name—The name of the Monitor that you have specified at the time of creating the Monitor.
 - Monitor description—The description of the Monitor that you have entered at the time of creating the Monitor.
 - Last Started—The date and time when the Monitor was last started.
 - Monitor status—The overall status of the Monitor that you run.
- View Monitor results for a specific period—You can select a predefined period (15 m, 2h, 4h, 8h, 16h, 24h, 1w) for which you want to view the results of all Monitors. You can also view the last started time of the Monitor.

Click **Custom** to set a custom time range for which you want to view the results of all Monitors. On the **Custom Time Range Selection** page that appears, enter the day and time in the **From** and **To** fields, respectively.

- View the list of events—You can view the events generated by the Monitor, and the date and time at which the event has occurred. The Event bar displays events based on the severity levels. The high-priority event, **Critical**, is displayed at the top of the events list, while low-priority event, **Information**, is displayed at the end of the list.

To view the list of all the events generated in the order of occurrence, click the **More** option. When you click **More**, the Events page appears. You can also sort the columns in the Events page. For more information on the events generated by the Monitor, see [Table 310 on page 1255](#).

- View the events on an event bar—You can view all the events generated on an event bar. The event bar is a color-driven bar graph that indicates events and the time at which the event has occurred. The colors represent the severity of events. That is, critical events are represented in red, errors are represented in orange, warnings are represented in yellow, information is represented in blue.

You can also hover over an event generated on the event bar to view the start time, end time, total number of events generated, and the number of events generated in each type.

- View the status of each Task within the Step—A Step is a collection of Tasks and a Monitor can only have one Step. Each Task within a Step is displayed as different event bars. Click > to view the event bar for each Task. When you hover over the event bar, you can view the events generated for the Task.
- View the list of Streams—You can view the list of Streams generated for the Task you selected and their details as displayed in [Table 311 on page 1256](#). To view the details of a specific Stream, click a *Stream-Name*. The *Stream-Name* page displays the result for metrics as area-graphs. For more information, see ["View Stream Details" on page 1258](#).
- Group all the Tasks—By default, all Streams are displayed in a table format. When you enable the **Group by Task** toggle button, the Streams are categorized based on the Tasks. By default, this toggle button is disabled.

When you enable this toggle button, you can expand and hide the Streams within the Tasks by using **Expand** and **Minimize** icons. When you click **Expand** icon, the Stream details for each Task are displayed, and when you click **Minimize**, the Stream details for each Task are hidden.

- Start or Stop a Monitor—You can start or stop a Monitor by using the **Start** or **Stop** options. If the Monitor status is Completed, click **More** on the top-right corner of the page and select **Start** to restart a Monitor. If the Monitor status is Running, click **More** on the top-right corner of the page and select **Stop** to stop a Monitor.

A message confirms that Routing Director successfully started or stopped the Monitor and the overall status of the Monitor changes to Running or Stopped.

- Edit a Monitor—You can edit a Monitor by using the Edit option.

To edit a Monitor, click **More** on the top-right corner of the page and select **Edit**.

You are redirected to the Edit Monitor (**Observability > Active Assurance > Measurement Designer**) page and the Monitor you configured is displayed in edit mode.

- Clone a Monitor—You can clone a Monitor by using the Clone option. To clone a Monitor, click **More** and select **Clone**. When you clone, an editable copy of the Monitor you configured is displayed.

You are redirected to the Measurement Designer (**Observability > Active Assurance > Measurement Designers**) page and a clone of the Monitor you configured is displayed.

- Delete a Monitor—You can delete a Monitor by using the Delete option.

To delete a Monitor, click **More** on the top-right corner of the page and select **Delete**. A message is displayed asking for your confirmation. Click **Yes**. A message confirms that Routing Director has successfully deleted the Monitor and you are redirected to the Monitor (**Observability > Active Assurance > Monitors**) page.

Table 310: Fields in the Events page

Field	Description
Severity	The type of severity level of the event that is raised when there is a criteria violation. The following are the severity levels: Critical—Indicates that the event is critical and needs immediate attention. Warning—Indicates that the event needs to be fixed but does not require immediate attention. Info—Indicates that an event is raised that provides information on progress of the Task and does not require attention. Error—Indicates that the event needs to be fixed and requires immediate attention and troubleshooting.
Description	The description that you specified when you configured the evaluation criteria for a Monitor.

Table 310: Fields in the Events page *(Continued)*

Field	Description
Raise Time	The date and time when the event was generated. The timestamp is displayed in the following format: Month DD, YYYY, HH:MM:SS AM/PM. The date and time is displayed according to the Raise delay you specified when you configured the evaluation criteria for a Monitor. For example, Mar 5, 2024, 4:29:52 PM.
Clear Time	The date and time when the event was cleared. The timestamp is displayed in the following format: Month DD, YYYY, HH:MM:SS AM/PM. The date and time is displayed according to the clear delay you specified when you configured the evaluation criteria for a Monitor. For example, Mar 5, 2024, 4:29:52 PM.
Subject	The details of the event generated is displayed in JSON format. It displays various IDs related to the event like Test Agent ID, Monitor ID, Task ID, Stream ID, and so on. To view the details, click show hyperlink.
Data	The details of the evaluation criteria associated with the event generated. To view data, click show hyperlink.

Table 311: Fields in the Monitor Result Table

Field	Description
Stream Name	The name of the Stream that is auto-generated by Routing Director when you run a Monitor.

Table 311: Fields in the Monitor Result Table (*Continued*)

Field	Description
Config	The configuration details of the selected Stream. Click show to view the details of the results metrics measured by the Stream. You can also copy the configuration.
Task Name	The name of the Task that generated the Stream. Tasks such as UDP is bi-directional. A bi-directional Task produces two Streams.
Task Status	The status of the Task generated by the Stream. The Task can have Waiting, Started, Running, Error, Stopped, and Completed status.
Severity	The severity level of events raised by Streams. The following are the severity levels: Critical—Indicates that the event is critical and needs immediate attention. Warning—Indicates that the event needs to be fixed but does not require immediate attention. Info—Indicates that an informational event is raised that provides details on the progress of the Task and does not require attention. Error—Indicates that the event needs to be fixed and requires immediate attention and troubleshooting.
Event Overview	The event overview of a Stream displays an event bar. You can view the start time, end time, total number of events generated, and the number of events generated in each type. The events can have Critical (Red), Error (Orange), Warning (Yellow), or Info (Blue) status. If you click a <i>Stream-Name</i>, you can access a detailed result of all metrics measured by the Stream. See "View Stream Details" on page 1258 for more information.

RELATED DOCUMENTATION

[View Stream Details](#) | 1258

View Stream Details

IN THIS SECTION

- [Streams Overview](#) | 1258
- [Access *Measurement-Name* Details Page](#) | 1258
- [Tasks You Can Perform](#) | 1259
- [Analyze the Stream Graph](#) | 1287

Streams Overview

At the time of Test or Monitor creation, you can configure evaluation criteria for various metrics.

When you run a Test or a Monitor, Routing Director instructs the Test Agents to send or receive traffic in the network by using the selected protocols. Each measurement produces one or more Stream of metrics.

The Test and Monitor evaluate these Streams and present a summary in the respective *Measurement-Name* Details pages. On violation of a configured evaluation criterion, an event is generated. You can view the Stream graphs on *Measurement-Name* Details page and identify the cause of violation.

Access *Measurement-Name* Details Page

To access the *Stream-Name* Details page:

1. Select **Observability > Active Assurance > Tests** or **Observability > Active Assurance > Monitors**, and click a *Test-Name* or a *Monitor-Name*.

The *Test-Name* or *Monitor-Name* page appears.

2. Click a *Stream-Name* from the Streams table.

The *Stream-Name* Details page appears.

Tasks You Can Perform

On the *Measurement-Name* Details page, you can:

- View Stream results for a specific period—You can select a predefined period (15m, 2h, 4h, 8h, 16h, 24h, 1w) for which you want to view the results of all Streams. You can also click **Custom** to set a custom time range for which you want to view the results of all Streams. When the Custom Time Range selection page appears, specify the day and time in the **From** and **To** fields, respectively.
- View the list of events—You can view the events generated by the Streams, and the date and time at which the event has occurred. The Event bar displays events based on the severity levels. The high-priority event, **Critical**, is displayed at the top of the events list, while the low-priority event, **Information**, is displayed at the end of the list.

To view the list of all the events generated in the order of occurrence, click the **More** option. When you click **More**, the Events page appear. You can also sort the columns in the Events page. For more information on the events generated by the Streams, see [Table 312 on page 1260](#).

- View the events on an event bar—You can view all the events generated on an event bar. The event bar is a color-driven bar graph that indicates events and the time at which the event has occurred. The colors represent the severity of events. That is, critical events are represented in red, errors are represented in orange, warnings are represented in yellow, information is represented in blue.

You can also hover over an event generated on the event bar to view the start time, end time, total number of events generated, and the number of events generated in each type.

- View the status of the selected Stream—You can view the status of the selected Stream. The Stream is displayed as an event bar. Click > to view the event bar for the Stream. When you hover over the event bar, you can view the events generated for the Stream.
- View the configurations details—You can view all the parameters you specified for a Task when you created a Test or a Monitor. Click **Config** to view all the parameters you have configured. For more information, see ["Create a Test" on page 958](#) and ["Create a Monitor" on page 1110](#).
- View the stream graph for a specific metric—You can view a stream graph of a specific metric. To view a stream graph, enable the *Metric-Name* toggle button for which you want to view the stream graph. The *Metric-Name* can vary based on the Tasks you have selected. A list of all the metrics are defined in [Table 313 on page 1261](#).
- View stream graphs for all metrics—You can view the stream graphs of all the metrics. To display stream graphs of all the generated metrics, click the **Show All** button.
- Group stream graphs by metrics—You can organize the stream graphs logically to get a coherent view of metrics.

When you enable the **Group Metrics** toggle button, the metrics are logically categorized. That is, all the related metrics are grouped under a primary stream graph. For example, for a ping Task, metrics

such as minimum time response, maximum time response, and average time responses will be grouped under Response (primary metric).

Once enabled, you can select and deselect the individual metrics in an stream graph. The data related to the metrics you selected will be displayed on the stream graph.

When you disable the **Group Metrics** toggle button, all the metrics are displayed. By default, this toggle button is disabled.

- View stream graphs in compact view—You can view all the stream graphs in compact mode to navigate easily between individual graphs. The compact view provides a streamline view of multiple graphs. To view the stream graphs in compact mode, enable the **Compact view** toggle button.

All the stream graphs for the metrics you toggled on will be displayed in compact view. The compact view provides a concise display of all the stream graphs that helps you to analyze various metrics simultaneously.

Table 312: Fields in the Events page

Field	Description
Severity	The type of severity level of the event that is raised when there is a criteria violation. The following are the severity levels: Critical—Indicates that the event is critical and needs immediate attention. Warning—Indicates that the event needs to be fixed but does not require immediate attention. Info—Indicates that an event is raised that provides information on progress of the Task and does not require attention. Error—Indicates that the event needs to be fixed and requires immediate attention and troubleshooting.
Description	The description that you specified when you configured the evaluation criteria for a Monitor.

Table 312: Fields in the Events page *(Continued)*

Field	Description
Raise Time	The date and time when the event was generated. The timestamp is displayed in the following format: Month DD, YYYY, HH:MM:SS AM/PM. The date and time is displayed according to the Raise delay you specified when you configured the evaluation criteria for a Monitor. For example, Mar 5, 2024, 4:29:52 PM.
Clear Time	The date and time when the event was cleared. The timestamp is displayed in the following format: Month DD, YYYY, HH:MM:SS AM/PM. The date and time is displayed according to the clear delay you specified when you configured the evaluation criteria for a Monitor. For example, Mar 5, 2024, 4:29:52 PM.
Subject	The details of the event generated is displayed in JSON format. It displays various IDs related to the event like Test Agent ID, Monitor ID, Task ID, Stream ID, and so on. To view the details, click show hyperlink.
Data	The details of the evaluation criteria associated with the event generated. To view data, click show hyperlink.

Table 313: Stream View Metrics

Stream Metric	Description
DNS	

Table 313: Stream View Metrics *(Continued)*

Stream Metric	Description
Minimum response time	Minimum response time of all the requests in the result interval. Response time is the time elapsed from when the DNS request was sent until a response was received. The response time is calculated in milliseconds.
Average response time	Average response time of all the requests in the result interval. Response time is the time elapsed from when the DNS request was sent until a response was received. The response time is calculated in milliseconds.
Maximum response time	Maximum response time of all the requests in the result interval. Response time is the time elapsed from when the DNS request was sent until a response was received. The response time is calculated in milliseconds.
ES timeout	ES timeout is the errored-seconds raised when Test Agent considers a DNS request to be timed-out. The ES timeout is calculated in milliseconds.
ES lifetime	ES lifetime is the errored-seconds raised when Test Agent fails the DNS request as there is no response received within the specified Request lifetime. The ES lifetime is calculated in milliseconds.
ES response	ES response is the errored-seconds raised when the DNS Response code differs from the expected response code or when the response differs from the expected response.

Table 313: Stream View Metrics *(Continued)*

Stream Metric	Description
ES	ES is the number of seconds during which error second thresholds are violated. It is the total number of error-seconds raised during result interval. The ES is calculated in milliseconds.
HTTP	
Connect time	Connect time is the time taken to set up a TCP connection between Test Agent and the Web server by using a TCP handshake. The connect time is calculated in milliseconds.
First byte received	First byte received indicates the total time taken by Test Agent to receive the first byte of response packet from the Web server. The first byte received time is calculated in milliseconds.
Minimum response time	Minimum response time is the duration when the request was sent until the full response is received. TCP connection setup time is included in the average response duration. The response time is calculated in milliseconds.
Average response time	Average response time is the average duration when the request was sent until the full response is received. TCP connection setup time is included in the average response duration. The response time is calculated in milliseconds.

Table 313: Stream View Metrics *(Continued)*

Stream Metric	Description
Maximum response time	Maximum response time is the maximum duration when the request was sent until the full response is received. TCP connection setup time is included in the average response duration. The response time is calculated in milliseconds.
Size	Size indicates the total size of the data transferred from a remote endpoint to Test agent in response to an HTTP request. The size also includes the header size. The size is calculated kilobytes.
Rate	Rate is the download rate at which the response is received from first byte to last. The rate is displayed in Megabits Per Second.
ES timeout	ES timeout is the number of errored-seconds raised because no HTTP response was received before the timeout period expires. The ES timeout is calculated in milliseconds.
ES response	ES response is the errored-seconds raised when the response code differs from the expected response code or when the response content differs from the expected response.
ES	ES is the number of seconds during which error-second thresholds were violated. It is the total number of errored-seconds raised during the result interval. The ES is calculated in milliseconds.
Ping	

Table 313: Stream View Metrics *(Continued)*

Stream Metric	Description
Successful ping counts	Number of ICMP or UDP echo request which has received a response during the result interval. The higher ping counts denote to a more reliable the network connection.
Minimum round-trip delay	Minimum round-trip delay of all the requests in the result interval. Round-trip measures the duration from the time Test Agent sends a request until it receives a response from the device. The delay is calculated in milliseconds.
Maximum round-trip delay	Maximum round-trip delay of all the requests in the result interval. Round-trip measures the duration from the time Test Agent sends a request until it receives a response from the remote endpoint. The delay is calculated in milliseconds.
Average round-trip delay	Average round-trip delay of all the requests in the result interval. Round-trip measures the duration from the time Test Agent sends a request until it receives a response from the remote endpoint. The delay is calculated in milliseconds.
Average round-trip DV	Average round-trip delay variance in the result interval. Round-trip delay variance is the difference in time between the request with longest round-trip time and the shortest round-trip time. The delay variance is calculated in milliseconds.
Lost	Indicates the number of ping request packets sent by Test Agent that did not receive a response before the configured lifetime.

Table 313: Stream View Metrics (Continued)

Stream Metric	Description
Loss	Indicates the number of ping request packets sent by Test Agent that did not receive a response before the configured lifetime.
ES	ES is the number of seconds during which error-second thresholds were violated. It is the total number of errored-seconds raised during the result interval. The ES is calculated in milliseconds.
ES loss	ES loss is the number of errored-seconds during the result interval where a request or more did not receive a response. The ES loss is calculated in milliseconds.
ES delay	ES delay is the number of errored-seconds raised when Test Agents experience round-trip delay greater than the configured timeout. The ES delay is calculated in milliseconds.
ES response	ES response is the number errored-seconds raised because Test Agent failed to resolve the destination hostname or because the response payload does not match what was sent.
ES delay variance (DV)	ES delay variance is the number of errored-seconds raised because Test Agent measures a variance in delay greater than the configured delay variation threshold. The ES delay variance is calculated in milliseconds.
RPM CISCO TWAMP NOTE: The RPM CISCO TWAMP plug-in measurement is a beta feature in this release.	

Table 313: Stream View Metrics *(Continued)*

Stream Metric	Description
Keep alive count	Keep alive count indicates the number of active Cisco TWAMP plug-in measurement updates in real time.
IPTV MPEG	
Rate	Rate indicates the bit rate of the MPEG program stream. The rate is displayed in Megabits Per Second.
Transport rate	Transport rate indicates the bit rate of the MPEG transport stream (MPEG-TS). It includes the overhead from the header of the Transport Stream packet. The Transport rate is displayed in Megabits Per Second.
MPEG loss	MPEG loss indicates the MPEG packet loss during a transmission. The loss is calculated from the continuity count (CC) errors in the MPEG stream.
PCR jitter	Program Clock Reference (PCR) refers to the variation in the arrival time of PCR packets when compared to the expected arrival time. The PCR jitter is displayed in milliseconds.
PAT error	A PAT error occurs if a Program Allocation Table (PAT) is not received on a multicast group within the configured PAT/PMT interval.
PMT errors	A PMT error occurs if a Program Map Table (PMT) is not received on a multicast group within the configured PAT/PMT interval.
PID errors	A PID error occurs if no frame is received within the specified PID interval threshold for every second that elapses.

Table 313: Stream View Metrics *(Continued)*

Stream Metric	Description
RTP errors	Real-Time protocol (RTP) error indicates the number of lost packets in the RTP header. Whether an MPEG stream contains RTP headers depends on the configuration of the encoder at the head-end.
RTP misordering	RTP misorders indicates the number of misordered packets based on Real-Time protocol (RTP) header. A packet is counted as misordered if it is delivered after a packet with a higher sequence number. Whether an MPEG stream contains RTP headers depends on the configuration of the encoder at the head-end.
RTP jitter	Real-Time protocol jitter indicates the variation in the arrival time of RTP packets when compared to the expected arrival time. If the MPEG stream contains RTP header, RTP jitter is calculated by using the timestamps from the Real Time protocol (RTP). Whether an MPEG stream contains RTP headers depends on the configuration of the encoder at the head-end.
ES	ES is the number of seconds during which errors occurred in a specific duration of time. ES is an aggregate of all specific ES metrics. The ES is displayed in seconds.
ES MPEG Loss	ES MPEG loss is the errored-seconds raised due to MPEG loss packet loss (CC errors). It is raised when the MPEG loss exceeds the configured MPEG loss threshold.

Table 313: Stream View Metrics *(Continued)*

Stream Metric	Description
ES Jitter	ES jitter is the number of errored-seconds raised when the PCR jitter or RTP jitter exceeds the defined jitter threshold.
ES invalid Stream	ES invalid stream is an aggregate of PAT, PMT, and PID errors. If any of these types of errors are experienced in a second, an errored-second is raised.
Netflix Speedtest	
Bandwidth speed	Bandwidth speed is the speed of data transmission to the Netflix servers or from the Netflix servers. The Bandwidth speed is displayed in Megabits Per Second.
End to End Latency	End-to-end latency indicates the latency occurred during the transfer of data from the Test Agent to the Netflix server or vice versa. The End to End Latency is displayed in milliseconds.
ES bandwidth	ES bandwidth indicates the number of errored-seconds raised if the bandwidth drops below the defined value.
ES latency	ES latency indicates the number of errored-seconds raised if the latency exceeds the defined value.
ES	ES is the number of seconds during which errors occurred in a specific duration of time. ES is an aggregate of all specific ES metrics. The ES is displayed in seconds.
ES rate	ES Rate is the number of errored-seconds occurred when the rate at which the UDP packet reception drops below the defined value.

Table 313: Stream View Metrics *(Continued)*

Stream Metric	Description
OTT - HLS	
Playback rate	Playback rate indicates the actual data rate of the video stream. The rate is displayed in Megabits Per Second.
Download rate	Download rate indicates the download rate of the segments. When the amount of buffered data drops below the defined buffer size, the download begins. The rate is displayed in Megabits Per Second.
Selected rate	Selected rate indicates the data rate of the selected variant in the manifest file. The rate is displayed in Megabits Per Second.
Buffer length	Buffer length is the duration of the initial buffered data at a given time. When the duration of the buffered data drops below the defined value, new segments are downloaded.
ES	ES is the number of seconds during which errors occurred in a specific duration of time. ES is an aggregate of all specific ES metrics.
ES playback rate	ES playback rate is the number of errored-seconds raised when the playback rate drops below the defined playback rate value.
ES selected rate	ES selected rate is the number of errored-seconds raised when selected rate drops below the defined value.

Table 313: Stream View Metrics *(Continued)*

Stream Metric	Description
ES download rate	ES download rate is the number of errored-seconds raised when the download rate drops below the defined value.
ES buffer	ES buffer is the number of errored-seconds raised when buffer length of the video stream drops below the defined value.
ES buffering	ES buffer is the number of errored-seconds raised when buffer length of the video stream drops below the defined value.
TCP	
Rate	Rate indicates the speed or throughput (goodput) at which the Test Agent receives the TCP packets in a selected interval of time. The rate is displayed in Megabits Per Second.
ES	ES is the number of seconds during which errors occurred in a specific duration of time. ES is an aggregate of all specific ES metrics. The ES is displayed in seconds.
ES rate	ES Rate is the number of errored-seconds occurred when the data rate drops below the defined value. The ES rate is displayed in seconds.
UDP	
Misordered packets	Misordered packets indicates the total number of packets received by the Test Agent in a different order than they were sent.

Table 313: Stream View Metrics *(Continued)*

Stream Metric	Description
Rate	Rate indicates the speed or throughput at which Test Agent receives the UDP packets in a selected interval. The rate is displayed in Megabits Per Second.
ES	ES is the number of seconds during which errors occurred in a specific duration of time. ES is an aggregate of all the specific ES metrics.
ES Rate	ES Rate is the number of errored-seconds occurred when the rate at which the UDP packet reception drops below the defined value. The rate is displayed in milliseconds.
ES Loss	ES loss is the number of errored-seconds raised when measured packet loss exceeds the defined loss value.
ES Delay	ES Delay is the number of errored-seconds raised when the measured one-way delay exceeds the defined delay value. The ES delay is displayed in milliseconds.
ES DSCP	ES DSCP is the number of errored-seconds occurred when the DSCP marked packets experience loss, delay, delay variance, or configuration issues.
SES	Severely errored second (SES) indicates the interval of time during which any of the defined SES thresholds were violated.
UAS	The Unavailable Seconds (UAS) metric indicates the number of seconds at which the service can be considered to be unavailable.

Table 313: Stream View Metrics *(Continued)*

Stream Metric	Description
Average delay	Average delay indicates the average delay experienced by data packets in the UDP measurement. The average delay is displayed in milliseconds.
Minimum delay	Minimum delay indicates the minimum delay experienced by data packets in a UDP measurement. The minimum delay is displayed in milliseconds.
Maximum delay	Average delay indicates the average delay experienced by data packets in the UDP measurement. The maximum delay is displayed in milliseconds.
Jitter	Jitter is the difference between maximum and minimum one-way delay measured during a result interval. The jitter is displayed in milliseconds.
Number of packets in the interval	Number of packets in the interval.
Percentage of packets lost	Loss percentage indicates the percentage of packets sent from the Test Agent that were lost before reaching the receiving Test Agent.
Lost Packet count	Lost indicates the number of packets sent from the Test Agent that were lost before reaching the receiving Test Agent.
TWAMP/TWAMP Light	

Table 313: Stream View Metrics *(Continued)*

Stream Metric	Description
Rate	Rate indicates the rate at which sender Test Agent receives response packets from the TWAMP reflector in a selected interval of time. The rate is calculated as the size of the response received divided by the total response time. The rate is calculated in Megabits Per Second.
Minimum round trip time	Minimum round-trip delay of all the requests in the result interval. Round-trip time indicates the duration from when a test packet is sent from a Test Agent until the packet is reflected back and received again on the Test Agent. Time spent in the reflector is not included. The time is calculated in milliseconds.
Average round trip time	Average round-trip delay of all the requests in the result interval. Round-trip time indicates the duration from when a test packet is sent from a Test Agent until the packet is reflected and received again on the Test Agent. Time spent in the reflector is not included. The time is calculated in milliseconds.
Maximum round trip time	Maximum round-trip delay of all the requests in the result interval. Round-trip time indicates the duration from when the test packet is sent from a Test Agent until the packet is reflected and received again on the Test Agent. Time spent in the reflector is not included. The time is calculated in milliseconds.
Round trip time DV	Average round-trip delay variance in result interval. Round-trip delay variance is the difference in round-trip time between the packets with the longest and the shortest round-trip time. The Round trip time DV is calculated in milliseconds.

Table 313: Stream View Metrics *(Continued)*

Stream Metric	Description
Received Packets	Received packets indicates the total number of packets received by Test Agent from the reflector.
Far-end loss	Far-end loss indicates the percentage of packets lost between Test Agent and the Reflector. It is the percentage of packets that were lost before reaching the reflector.
Far-end lost	Far-end lost indicates the number of packets lost before reaching the reflector.
Far-end misorders	Far-end misorders indicate the number of packets sent by Test Agent that reached the reflector out of sequence compared to their original order of transmission.
Minimum far-end delay	Minimum one-way far-end delay indicates the minimum delay measured for a packet to travel from Test Agent to the Reflector. The delay is calculated in milliseconds.
Average far-end delay	Average one-way far-end delay indicates the average delay measured for a packet to travel from Test Agent to the Reflector. The delay is calculated in milliseconds.
Maximum far-end delay	Maximum one-way far-end delay indicates the average delay measured for a packet to travel from Test Agent to the Reflector. The delay is calculated in milliseconds.

Table 313: Stream View Metrics *(Continued)*

Stream Metric	Description
Far-end DV	Far-end DV indicates the delay variance measured at the remote endpoint. It is the difference between the maximum and minimum one-way far-end delay. The delay is calculated in milliseconds.
Near-end loss	Near-end loss indicates the percentage of packets lost between the Reflector and Test Agent. It is the percentage of packets that were lost before reaching the Test Agent.
Near-end lost	Near-end lost indicates the number of packets lost between the Reflector and Test Agent. It is the number of packets that were lost before reaching the Test Agent.
Near-end misorders	Near-end misorders indicates the number of packets that reached Test Agent out of sequence compared to their original order of transmission.
Minimum near-end delay	Minimum one-way near-end delay indicates the minimum delay experienced by a packet to travel from the Reflector to Test Agent. The delay is calculated in milliseconds.
Average near-end delay	Average one-way near-end delay indicates the average delay experienced by a packet to travel from the Reflector to Test Agent. The delay is calculated in milliseconds.
Maximum near-end delay	Maximum one-way near-end delay indicates the maximum delay experienced by a packet to travel from the Reflector to Test Agent. The delay is calculated in milliseconds.

Table 313: Stream View Metrics *(Continued)*

Stream Metric	Description
Near-end DV	Near-end DV indicates the delay variance measured at Test Agent endpoint. It is the difference between the maximum and minimum one-way near-end delay. The delay is calculated in milliseconds.
ES	ES is the number of seconds during which error-second thresholds were violated. It is the total number of errored-seconds raised. The ES is calculated in milliseconds.
ES delay	ES delay is the number of errored-seconds raised when Test Agent measures a delay greater than the configured delay threshold. The ES delay is calculated in milliseconds.
ES loss	ES loss is the number of errored-seconds raised when the measured loss is greater than the configured loss threshold. The ES loss is calculated in milliseconds.
ES delay variance (DV)	ES delay variance is the number of errored-seconds raised when the measured delay variance is greater than the configured DV threshold. The ES delay variance is calculated in milliseconds.
ES DSCP	ES DSCP is the number of errored-seconds occurred when the received DSCP value does not match the expected value.
SES	Severely errored second (SES) indicates the number of seconds that a SES threshold has been violated.

Table 313: Stream View Metrics *(Continued)*

Stream Metric	Description
First round trip delay percentile	First round trip delay percentile. If the delay exceeds the configured value, the packet is included in the defined first delay percentile slot.
Second round trip delay percentile	Second round trip delay percentile. If the delay exceeds the configured value, the packet is included in the defined second delay percentile slot.
First far end round trip delay percentile	First far end round trip delay percentile. If the delay exceeds the configured value, the packet is included in the defined first far end round trip delay percentile slot.
Second far end round trip delay percentile	Second far end round trip delay percentile. If the delay exceeds the configured value, the packet is included in the defined second far end round trip delay percentile slot.
First near end round trip delay percentile	First near end round trip delay percentile. If the delay exceeds the configured value, the packet is included in the defined first near end round trip delay percentile slot.
Second near end round trip delay percentile	Second near end round trip delay percentile. If the delay exceeds the configured value, the packet is included in the defined second near end round trip delay percentile slot.
ES for first round trip delay percentile	ES for first round trip delay percentile. If the delay exceeds the configured value, the packet is included in the defined first round trip delay percentile slot and an errored-second is raised.
ES for second round trip delay percentile	ES for second round trip delay percentile. If the delay exceeds the configured value, the packet is included in the defined second round trip delay percentile slot and an errored-second is raised.

Table 313: Stream View Metrics *(Continued)*

Stream Metric	Description
Round trip loss	Round-trip loss indicates the percentage of TWAMP packets sent from Test Agent that were lost either on the way to the reflector or on the way back to Test Agent. The loss percentage is calculated by comparing the total packets lost with the total number of packets sent.
Round trip lost	Round-trip lost indicates the number of TWAMP packets sent from Test Agent that were lost either on the way to the reflector or on the way back to Test Agent.
Timestamp samples	The number of valid timestamp samples used in delay measurements.
TWAMP Reflector	
Received packets	Received packets indicates the number of packets received by the Test Agent from the reflector.
Rate	Rate indicates the rate at which Test Agent received TWAMP traffic. It is the aggregate of all the concurrent sessions. The rate is calculated as the size of the response received divided by the total response time. The rate is displayed in Megabits Per Second.
Active sessions	Active sessions indicate the number of currently active sessions on the specified reflector.
Bytes received	Bytes received indicates the total size of TWAMP traffic received during the report interval. The size also includes the header size.

Table 313: Stream View Metrics *(Continued)*

Stream Metric	Description
ES	ES is the number of seconds during which errors occurred in a specific duration of time. It is the total number of error-seconds that any threshold was violated. The ES is calculated in milliseconds.
ES rate	ES Rate is the number of errored-seconds raised because the rate threshold was violated. The ES rate is calculated in milliseconds.
RPM (HTTP, PING, TCP, UDP)	
Round trip time	Round-trip time (RTT) measures the amount of time taken by the request packet to travel from Test Agent Devices to the remote endpoint and the response packet to travel from the remote endpoint to Test Agent Devices. RTT indicates the time delay between sending a request and receiving a response. RTT indicates the time delay between sending a request and receiving a response.
Round trip jitter	Jitter occurs when the packets experience a delay during a round trip, from the Test Agent to the remote endpoint. The delay is calculated in milliseconds.
Round trip inter-arrival	The Round-trip interarrival jitter indicates an estimate of the total statistical variance of a packet's interarrival time as defined in IETF RFC 1889. The interarrival jitter is calculated in milliseconds

Table 313: Stream View Metrics *(Continued)*

Stream Metric	Description
Loss	Loss indicates the percentage of request packets sent for which responses were not received.
Y.1731 DM	
Minimum two-way frame delay	Minimum two-way frame delay indicates the minimum delay for each frame to travel from local maintenance end point (MEP) to the remote maintenance end point (MEP) and back to the local MEP.
Maximum two-way frame delay	Maximum two-way frame delay indicates the maximum delay for each frame to travel from local MEP to the remote MEP and back to local MEP.
Average two-way frame delay	Average two-way frame delay indicates the average delay for each frame to travel from local MEP to the remote MEP and back to the local MEP.
Minimum one-way frame delay at far end	Minimum one-way frame delay at far end indicates the minimum delay for each frame to travel from local MEP to the remote MEP.
Maximum one-way frame delay at far end	Maximum one-way frame delay at far end indicates the maximum delay for each frame to travel from local MEP to the remote MEP.
Average one-way frame delay at far end	Average one-way frame delay at far end indicates the average delay for each frame to travel from local MEP to the remote MEP.
Minimum one-way frame delay at near end	Indicates the minimum delay for each frame to travel back from the remote MEP to local MEP.

Table 313: Stream View Metrics *(Continued)*

Stream Metric	Description
Maximum one-way frame delay at near end	Minimum one-way frame delay at near end indicates the maximum delay for each frame to travel back from the remote MEP to local MEP.
Average one-way frame delay at near end	Average one-way frame delay at near end indicates the average delay for each frame to travel back from the remote MEP to local MEP.
Minimum two-way inter frame delay variance	Minimum two-way inter frame delay variance indicates the minimum delay variance experienced to travel from local MEP to the remote MEP and back to the local MEP.
Maximum two-way inter frame delay variance	Maximum two-way inter frame delay variance indicates the maximum delay variance experienced to travel from local MEP to the remote MEP and back to the local MEP.
Average two-way inter frame delay variance	Average two-way inter frame delay variance indicates the average delay variance experienced to travel from local MEP to the remote MEP and back to the local MEP.
Minimum delay variance at far end	Minimum delay variance at far end indicates the minimum delay variance experienced by each frame to travel from local MEP to the remote MEP.
Maximum frame delay variance at far end	Maximum inter frame delay variance at far end indicates the minimum delay variance experienced by each frame to travel from local MEP to the remote MEP.
Average frame delay variance at far end	Average frame delay variance at far end indicates the average delay variance experienced by each frame to travel from local MEP to the remote MEP.

Table 313: Stream View Metrics *(Continued)*

Stream Metric	Description
Minimum frame delay variance at near end	Minimum frame delay variance at near end indicates the minimum delay variance experienced by each frame to travel back from the remote MEP to local MEP.
Maximum frame delay variance at near end	Maximum frame delay variance at near end indicates the maximum delay variance experienced by each frame to travel back from the remote MEP to local MEP.
Average frame delay variance at near end	Average frame delay variance at near end indicates the average delay variance experienced by each frame to travel back from the remote MEP to local MEP.
SOAM frames transmitted	SOAM frames received indicates the SOAM delay measurement messages transmitted by local MEP to the remote MEP to measure network performance.
SOAM frames received	SOAM frames received indicates the SOAM delay measurement messages received by the remote MEP from the local MEP to measure network performance.
Y.1731 SLM	
SOAM frames transmitted	SOAM frames transmitted indicates the SOAM synthetic loss measurement frames transmitted by local MEP to the remote MEP to monitor network performance.
SOAM frames received	SOAM frames received indicates the SOAM synthetic loss measurement frames received by the remote MEP from local MEP to monitor network performance.
Frames transmitted at far end	Frames transmitted at far end indicates the frames transmitted by local MEP to the remote MEP to monitor network performance.

Table 313: Stream View Metrics *(Continued)*

Stream Metric	Description
Frames received at far end	Frames received at far end indicates the frames received by local MEP to the remote MEP to monitor network performance.
Frames transmitted at near end	Frames transmitted at far end indicates the frames transmitted back by the remote MEP to local MEP to monitor network performance.
Frames received at near end	Frames received at far end indicates the frames transmitted back by the remote MEP to local MEP to monitor network performance.
Frame loss at far end	Frame loss at far end indicates the number of frames that were sent by local MEP but failed to receive at the remote MEP.
Frame loss at near end	Frame loss at near end indicates the number of frames that were sent back by the remote MEP but failed to receive at the local MEP.
Min frame loss ratio at far end	Minimum frame loss ratio indicates the minimum ratio of frames lost during transmission of frames from local MEP to the remote MEP.
Max frame loss ratio at far end	Maximum frame loss ratio indicates the maximum ratio of frames lost during transmission of frames from local MEP to the remote MEP.
Avg frame loss ratio at far end	Average frame loss ratio indicates the average ratio of frames lost during transmission of frames from local MEP to the remote MEP.
Min frame loss ratio at near end	Minimum frame loss ratio at near end indicates the minimum ratio of frames lost during transmission of frames back to local MEP from the remote MEP.

Table 313: Stream View Metrics *(Continued)*

Stream Metric	Description
Max frame loss ratio at near end	Maximum frame loss ratio at near end indicates the maximum ratio of frames lost during transmission of frames back to local MEP from the remote MEP.
Avg frame loss ratio at near end	Average frame loss ratio at near end indicates the average ratio of frames lost during transmission of frames back to local MEP from the remote MEP.
High loss intervals at far end	High loss intervals at far end indicates the intervals when the number of lost frames traveling from the local MEP to the remote MEP exceeds the configured threshold.
High loss intervals at near end	High loss intervals at near end indicates the intervals when the number of frames lost while traveling from the remote MEP to local MEP exceeds the configured threshold.
Consecutive high loss intervals at far end	Consecutive high loss intervals at far end indicates the consecutive intervals when the number of lost frames traveling from the local MEP to the remote MEP exceeds the configured threshold.
Consecutive high loss intervals at near end	Consecutive high loss intervals at near end indicates the consecutive intervals when the number of lost frames traveling from the remote MEP to local MEP exceeds the configured threshold.
Availability indicator at far end	Availability indicator at far end indicates the number of shorter intervals (deltas) within the measurement period that were lost when traveling from the local MEP to the remote MEP is below the configured threshold.

Table 313: Stream View Metrics *(Continued)*

Stream Metric	Description
Availability indicator at near end	Availability indicator at near end indicates the number of shorter intervals (deltas) within the measurement period that were lost when traveling from the remote MEP to local MEP is below the configured threshold.
Unavailability indicator at far end	Unavailability indicator at far end indicates the number of shorter intervals (deltas) within the measurement period that were lost when traveling from local MEP to the remote MEP exceeds the configured threshold.
Unavailability indicator at near end	Unavailability indicator at near end indicates the number of shorter intervals (deltas) within the measurement period that were lost when traveling from the remote MEP to local MEP exceeds the configured threshold.
Availability Frame Loss Ratio min at far end	Availability Frame Loss Ratio min at far end indicates the minimum ratio of frames lost to the total transmitted frames within the measurement period while traveling from local MEP to the remote MEP is below the configured threshold.
Availability Frame Loss Ratio max at far end	Availability Frame Loss Ratio max at far end indicates the maximum ratio of frames lost to the total transmitted frames within the measurement period while traveling from local MEP to the remote MEP is below the configured threshold.
Availability Frame Loss Ratio avg at far end	Availability Frame Loss Ratio avg at far end indicates the average ratio of frames lost to the total transmitted frames within the measurement period while traveling from local MEP to the remote MEP is below the configured threshold.


Analyze the Stream Graph

A stream graph is a graphical representation of a metric generated by a Stream over a period of time. You can generate multiple stream graphs for various metrics by enabling the corresponding toggle button.

In the stream graph, the **X** axis (horizontal) always represents the time duration of Test or Monitor. The **Y** axis (vertical) represents the metric value that is being measured.

For example, create a Test by adding an RPM Ping Task with the threshold value for Round Trip Time (RTT) set to 10 seconds (s). When you run this Test, the following graph is displayed in the *Stream-Name* Details page.

Figure 50: Sample Stream graph



In this stream graph ([Figure 50 on page 1287](#)), the **X** axis represents the time duration of the Test (14:15:10 through 14:17:10) and the **Y** axis displays the RTT value (0 through 14 milliseconds).

From the stream graph, you can infer that the RTT has exceeded the threshold (> 10) value at 14:16:16. At this moment, because the RTT value reached 11.468 ms, which is greater than the specified threshold value, an event is generated.

RELATED DOCUMENTATION

[Create a Test](#) | 958

[Create a Monitor](#) | 1110

Test Templates

IN THIS CHAPTER

- [About the Test Templates Page | 1288](#)
- [Create a Test Template | 1291](#)

About the Test Templates Page

IN THIS SECTION

- [Tasks You Can Perform | 1288](#)

To access this page, click **Inventory > Active Assurance > Test Templates**.

Test templates enable you to create customized Test templates and reuse these templates while creating Tests. Test templates eliminate the need to manually configure Task parameters and metrics evaluation criteria each time you create a Test. Instead, you can select an existing Test template and run a Test.

You (superusers or network administrators) can use the Test Templates page to view a list of Test templates in your organization. You can also create, clone, and delete Test templates. In addition, you can edit the name, tags, and description of a Test template.

Tasks You Can Perform

You can perform the following tasks on the Test Templates page:

- View the list of Test templates—You can view the list of all the Test templates in your organization. [Table 314 on page 1290](#) describes the fields on the Test Templates page.

- View the details of a Test template—You can select and view the details of a Test template. You can either click **More > Detail** or hover over the template and click the **Detailed View** icon. The Test Template Details pane appears on the right side of the page displaying the details.

Click the **Close (x)** icon to close the pane.

You can also click a *Test-Template-Name* to view the selected template on the Measurement Designer page. From the Measurement Designer page, you can edit the Test template details and save it as a new template.

- Create a Test template—You can create a Test template as described in ["Create a Test Template" on page 1291](#).
- Clone a Test template—To clone a Test template:

1. Select the *Test-Template* that you want to clone.
2. Click the **Clone** button on the top-right corner of the page.

You are redirected to the Edit Test (**Observability > Active Assurance > Measurement Designer**) page where a new template is displayed in the edit mode. The template name is displayed in Clone *Test-Template-Name* format.

3. Edit the template as desired, and click **Save**.

A message is displayed confirming that the Test template is saved successfully, and the Test template you have cloned is listed on the Test Templates (**Inventory > Active Assurance > Test Templates**) page.

- Edit a Test template details—To edit a Test template:
1. Click the **Edit** (pencil) icon on the top-right corner of the page.

The Test Template Details pane appears on the right side of the page displaying the Test templates overview details.

2. Edit the Name, Description, and Tags by using the **Edit** (pencil) icon.
3. Click **Save**.

Once you save, the updated details are displayed on the Test templates overview details pane.

- Delete a Test template— To delete the Test template:

1. Click **Delete** (trashcan) icon on the top-right corner of the page.

A message is displayed asking for your confirmation to delete the template.

2. Click **Yes**.

A message confirms that Routing Director has successfully deleted the Test template and you are redirected to the Test Templates (**Observability > Active Assurance > Tests**) page.

- You can also perform the following tasks on this page:
 - Sort, resize, or re-arrange columns in a table (grid).



NOTE: Sort functionality does not work for Created By, Updated By, and Tags.

- Show or hide columns in the table or reset page preferences, using the vertical ellipsis menu.
- Filter the data displayed in the table—Click the filter icon (funnel) and select whether you want to show or hide advanced filters. You can then add or remove filter criteria, save criteria as a filter, apply or clear filters, and so on. The filtered results are displayed on the same page.

For more information, see ["GUI Overview" on page 6](#).

Table 314: Fields on the Test Templates Page

Fields	Description
Name	Displays the name of the Test template that you specified when you created the Test template.
Description	Displays the description of the Test template that you specified when you created the Test template.
Tags	Displays the tags you have configured for the Test template in the <i>key:value</i> format. A tag is a key-value pair in which the key signifies a category for which you configure a value. The value is an identifier for the category. Examples of key-value pairs are device and device name (edgedevice:acx7000), site and site name (site:bengaluru).
Created	Displays the date and time when the Test template was created. The timestamp is displayed in the following format: Month DD, YYYY, HH:MM:SS AM/PM. For example, Jun 16, 2024, 4:29:52 PM.

Table 314: Fields on the Test Templates Page *(Continued)*

Fields	Description
Last Updated	Displays the latest date and time when the Test template was last updated. The timestamp is displayed in the following format: Month DD, YYYY, HH:MM:SS AM/PM. For example, Jun 16, 2024, 4:29:52 PM.
Created By	Displays the username (e-mail address) of the user who has created the Test template.
Updated By	Displays the username (e-mail address) of the user who has updated the Test template.

Create a Test Template

Test templates enable you to create customized Test templates and reuse these templates while creating Tests. Test templates eliminate the need to manually configure Task parameters and metrics evaluation criteria each time you create a Test. Instead, you can select an existing Test template and run a Test.

To create a Test template:

1. Do one of the following:
 - a. Navigate to the Measurement Designers (**Observability > Active Assurance > Measurement Designers**) page, or,
 - b. Navigate to the Test Templates (**Inventory > Active Assurance > Test Templates**) page and click the Add (+) icon.

You are redirected to the Measurement Designer page.

2. Specify a name for the Test template.

Click the **Sequence Name** or click the **Edit** (pencil) icon to enter a name for the Test template in the **Sequence Name** text box.

Sequence Name is the placeholder name for the Test template.



NOTE: Fields marked with an asterisk (*) are mandatory.

3. Add Steps and Tasks. For information on adding steps and tasks for a Test, see ["Create a Test" on page 959](#).
4. Click the **Save as Template** button on the top-right corner of the page.

A Save as Test Template page appears asking to confirm the save operation. You can also edit the name, description, and tags from this page.



NOTE: When you edit the name, specify a unique Test template name. The template does not get saved if you use a template name twice.

5. Click **Save**.

A message is displayed confirming that the Test template is saved successfully, and the Test template you have created is listed on the Test Templates (**Inventory > Active Assurance > Test Templates**) page.

RELATED DOCUMENTATION

| [Active Assurance Workflow](#) | 891

8

PART

Network Optimization

- Introduction | 1294
 - Network Optimization Intents | 1296
-

Introduction

IN THIS CHAPTER

- [Network Optimization Overview | 1294](#)

Network Optimization Overview

IN THIS SECTION

- [Benefits of Network Optimization | 1295](#)

Many modern networks are highly complex with numerous interconnected devices, making it challenging to identify and resolve optimization issues. In addition, with ever-increasing demand for bandwidth and with emerging technology, optimizing the dynamic traffic can be more challenging.

Network Optimization use case in Routing Director enables you to optimize the utilization of network resources, enhance network performance, and ensure reliable and efficient delivery of data across the network. This is achieved by managing the life-cycle of label-switched paths (LSPs), also known as tunnels or segment routing policies, through an intent-based approach.

You create a path intent using the Routing Director GUI. Path intents provide you a way to define the connectivity you want to achieve and define how the controller should optimize the connectivity based on the real-time network conditions. Unlike traditional methods where each tunnel must be provisioned individually with all its attributes configured from scratch, path intent allows you to create and manage profiles that can be reused. This modular approach reduces redundant configurations and streamlines the process of provisioning connectivity using multiple tunnels.

The path intent consists of the following three sub-profiles:

- Tunnel profile—Defines specific properties of the LSP or segment routing policy.
- Optimization profile—Defines when the LSP or segment routing policy must be recomputed.

- Endpoints profile—Defines the endpoints and the topology of the connectivity between those endpoints.

When you apply the intent to the network, Routing Director interprets these intent-based profiles, automates the creation, modification, and deletion of LSPs or segment routing policies, by autonomously executing the necessary actions to align the network state with the specified intent. Routing Director ensures that LSPs or segment routing policies are established based on network policies and traffic engineering constraints to achieve the desired service level agreements (SLAs).

The Network Optimization use case relies on the network telemetry to have visibility and insight on link bandwidth, tunnel bandwidth utilization, and measured link delay across the network, and dynamically adjusts LSP paths to meet changing traffic patterns and demands.

Benefits of Network Optimization

The intent-based approach eliminates the complexity of manually configuring the device and reduces the time taken to optimize your network. In addition, the intent-based approach simplifies network operations by eliminating human error, thereby enhancing operational efficiency and reliability.

RELATED DOCUMENTATION

[Path Intent Workflow](#) | 1296

[About the Network Optimization Intent Page](#) | 1298

Network Optimization Intents

IN THIS CHAPTER

- Path Intent Workflow | 1296
- About the Network Optimization Intent Page | 1298
- About the Path Intent Tab | 1300
- Add a Path Intent | 1302
- About the Tunnel Profiles Tab | 1304
- Add a Tunnel Profile | 1305
- About the Optimization Profiles | 1310
- Add an Optimization Profile | 1311
- About the Endpoint Profiles | 1313
- Add an Endpoint Profile | 1314
- About the Device Tab (Network Optimization) | 1316
- About the Link Tab (Network Optimization) | 1320
- About the Tunnels Tab (Network Optimization) | 1324

Path Intent Workflow

Traditionally, tunnels are provisioned on a per-instance basis directly from the network information. Path intent abstracts this process by breaking down the attributes of a tunnel into reusable components, known as sub-profiles. These sub-profiles include the tunnel profile, optimization profile, and endpoint profile. This modular approach allows you to define and manage these components separately and then combine them to create one or more tunnels. Path intents simplify tunnel management and enhances the efficiency.

The following procedure lists the step to create and apply a path intent to your network:

1. Login in to the Routing Director GUI. See ["Administration Workflow" on page 55](#).

2. Onboard devices. For information on onboarding the devices, see ["Device Onboarding Workflow" on page 173](#).



NOTE: When you onboard devices, ensure that you do the following while adding a device profile:

- Enable the **Traffic-Engineering** check box in the **Analytics** tab. You must do this for all devices.
- Add the PCEP VIP address in the PCEP settings of the **Protocols** tab.

3. Enable dynamic topology and ensure PCEP session is established with PCC. For more information, see ["Dynamic Topology Workflow" on page 555](#).

Dynamic topology requires configuration of BGP-LS settings on one router in the domain. See the following example:

```
set policy-options policy-statement pso-accept-link-state from family traffic-engineering
set policy-options policy-statement pso-accept-link-state then accept
set protocols bgp group BGP-LS type internal
set protocols bgp group BGP-LS local-address <Local-IP-Adress>
set protocols bgp group BGP-LS family traffic-engineering unicast
set protocols bgp group BGP-LS export pso-accept-link-state
set protocols bgp group BGP-LS peer-as 65200
set protocols bgp group BGP-LS allow 0.0.0.0/0
```

4. Navigate to **Network Optimization > Network > Network Optimization Intent**. On this page:
 - a. (Optional) On the Topology map, view the devices that you have onboarded and the links associated to those devices.
 - b. Navigate to the Tunnel Profiles tab and click + to create a tunnel profile. Ensure that you publish the tunnel profile. For more information, see ["Add a Tunnel Profile" on page 1305](#).
 - c. Navigate to the Optimization Profiles tab and click + to create an optimization profile. Ensure that you publish the optimization profile. For more information, see ["Add an Optimization Profile" on page 1311](#).
 - d. Navigate to the Endpoint Profiles tab and click + to create an endpoint profile. Ensure that you publish the endpoint profile. For more information, see ["Add an Endpoint Profile" on page 1314](#).
5. Navigate to the Path Intent tab and click + to create a path intent. For more information, see ["Add a Path Intent" on page 1302](#).

When you click the **Publish** button, the path intent is applied to the network. The tunnels are created and all the attributes and constraints specified in the profiles are applied consistently across the tunnels.

You can view and manage the tunnels on the network topology map. The topology map shows the start and end points of each tunnel, the path it takes, and any relevant attributes such as utilization and delay. You can filter and select tunnels to view detailed information and metrics, helping them monitor and manage the network effectively. This visualization is critical in understanding network topology and ensuring that tunnels meet the desired performance and reliability criteria.

RELATED DOCUMENTATION

[Network Optimization Overview](#) | 1294

About the Network Optimization Intent Page

To access the Network Optimization Intent page, click **Network Optimization > Network > Network Optimization Intent**.

The Network Optimization Intent page can be divided into three main components:

- **Interactive Topology Map**—Enables you to view devices, links, and tunnels across sites in your network and customize the map display. Manage the map view, set device positions on the map, customize device and link labels and font size using the right click menu.

For more information, see "[View Network Topology Details](#)" on page 504.

- **Topology Menu Bar**—A vertical bar at the top-right corner of the Network Optimization Intent page, which consists of the following:
 - Reset icon—Center the topology map so that it zooms to fit the screen.
 - Minus icon—Zoom out (reduce) the topology map.
 - Plus icon—Zoom in (enlarge) the topology map.
 - Switch to Cluster View/Node View icon—Switch to Cluster View from the default Node View.
 - Settings icon—Configure display properties of the topology map.

For more information, see "[View Network Topology Details](#)" on page 504

- **Network Table**—Displays detailed information about path intents, profiles (tunnel, optimization, and endpoint profiles), and network elements (devices, links, and tunnels).

Table 315 on page 1299 lists the tabs on the Network Optimization Intent page.

Click the collapsible arrow icon at the bottom-left of the topology map to display or hide the network table.

Table 315: Tabs on the Network Optimization Intent page.

Tab	Description
Path Intent	See "About the Path Intent Tab" on page 1300.
Tunnel Profiles	See "About the Tunnel Profiles Tab" on page 1304.
Optimization Profiles	See "About the Optimization Profiles" on page 1310.
Endpoint Profiles	See "About the Endpoint Profiles" on page 1313.
Device	See "About the Device Tab (Network Optimization)" on page 1316.
Link	See "About the Link Tab (Network Optimization)" on page 1320.
Tunnels	See "About the Tunnels Tab (Network Optimization)" on page 1324.

RELATED DOCUMENTATION

About the Path Intent Tab

IN THIS SECTION

- [Tasks You can Perform | 1300](#)
- [Field Descriptions | 1301](#)

To access the Path Intent tab, click **Network Optimization > Network > Network Optimization Intent**.

On this tab, you can create intents, view the status (Published or Draft) and the profiles (Tunnel, optimization, and endpoint profiles) linked to the intent.

Path intents implement the required LSP or segment routing policy configurations that define how traffic should be routed through the network. Unlike traditional methods where each tunnel must be provisioned individually with all its attributes configured from scratch, path intent allows you to create and manage sub-profiles that can be reused. This modular approach reduces redundant configurations and streamlines the process of provisioning multiple tunnels.

The path intent consists of the following three sub-profiles:

- Tunnel profile—Defines specific properties of the LSP.
- Optimization profile—Defines when the LSP must be recomputed.
- Endpoints profile—Defines the endpoints of the LSP.

For more information, see ["Network Optimization Overview" on page 1294](#) and ["About the Network Optimization Intent Page" on page 1298](#).

Tasks You can Perform

- View details of a path intent—You can select and view the details of a path intent. You can either click **More > Detail** or hover over the path intent and click the **Detailed View** icon. The Path Intent Details pane appears on the right side of the page displaying the details.

Click the **Close (x)** icon to close the pane.

- Create a path intent—Create a path intent to manage traffic flow in your network. Click **+** to create a path intent. For more information, see ["Add a Path Intent" on page 1302](#).
- Edit a path intent—To edit a path intent, select the path intent and click the pencil icon. After you edit, you can choose to save the path intent as a draft or publish the path intent.

- Delete a path intent—To delete a path intent, select a path intent and click the delete icon. A message confirms that Routing Director has successfully deleted the path intent.



NOTE: If you delete a path intent that is in published state, then the LSPs associated with the path intent are deleted.

To perform common actions such as filtering using advanced filter criteria and resetting preferences, see ["Network Table Overview" on page 515](#).

Field Descriptions

[Table 316 on page 1301](#) lists the fields on the Path Intent tab.

Table 316: Fields on the Path Intent Tab

Column	Description
Name	Displays the name of the path intent.
Description	Displays the description that you have entered at the time of path intent creation.
Tunnel Profile	Displays the name of the tunnel profile that you have linked to the path intent.
Optimization Profile	Displays the name of the optimization profile that you have linked to the path intent.
Endpoint Profile	Displays the name of the endpoint profile that you have linked to the path intent.
Status	Displays the status (Draft or Published) of the path intent.
Provisioned Tunnels	List the numbers of tunnels created for the path intent. Click the link to view the tunnel details.
Update Options	Displays information on whether the new tunnel is provisioned before deleting an old tunnel or an old tunnel is deleted before provisioning a new tunnel.

Add a Path Intent

Path intents implement specific LSP or segment routing policy configurations that define how traffic should be routed through the network. The path intent consists of the following three sub-profiles:

- Tunnel profile—Defines specific properties of the LSP.
- Optimization profile—Defines when the LSP must be recomputed.
- Endpoints profile—Defines the endpoints of the LSP.

For more information, see ["Network Optimization Overview" on page 1294](#) and ["About the Path Intent Tab" on page 1300](#).

To create a path intent profile:

1. Navigate to **Network Optimization > Network > Network Optimization Intent**.

The Topology page is displayed with the topology map at the center and the network information table at the bottom of the page.

2. Navigate to the Path Intent tab.

3. Click the plus icon (+).

The Create Path Intent page appears.

4. Complete the configuration according to the guidelines in [Table 317 on page 1303](#).



NOTE: Fields marked with an asterisk (*) are mandatory.

5. Do one of the following:

- Click **Save Draft**, to save the path intent. Save the path intent profile if you intend to modify it later.
- Click **Publish**, to publish the path intent. Publish the path intent if you want to apply the path intent to your network.

A message indicating that the path intent is created appears. You are returned to the Path Intent tab where you can view the newly added path intent.

Table 317: Fields on the Create Path Intent Page

Column	Description
Name*	Specify the name of the path intent. You cannot use space, comma, or semicolon.
Description	Enter a description for the path intent.
Tunnel Profile*	Select a tunnel profile from the list. A tunnel profile defines the properties of an LSP. Only published tunnel profiles are listed here.
Optimization Profiles*	Select an optimization profile from the list. An optimization profile defines the conditions when an LSP must be recomputed. Only published optimization profiles are listed here.
Endpoint group*	Select an endpoint group from the list. An endpoint group defines the source and destination endpoints of an LSP. Only published endpoint profiles are listed here.
Update Option*	Select whether you want a new tunnel to provisioned before deleting the old tunnel or you want an old tunnel to be deleted before provisioning the new tunnel.

RELATED DOCUMENTATION

[Path Intent Workflow](#) | 1296

About the Tunnel Profiles Tab

IN THIS SECTION

- [Tasks You can Perform | 1304](#)
- [Field Descriptions | 1305](#)

To access the Tunnel Profiles tab, click **Network Optimization > Network > Network Optimization Intent**.

On this tab, you can create tunnel profiles, view the status (Published or Draft) and the number of intents linked to the tunnel profile.

Unlike traditional methods where each tunnel must be provisioned individually with all its attributes configured from scratch, tunnel profiles define the specific attributes of the LSPs (tunnels) or segment routing policies, which can be later used to create path intents. This modular approach reduces redundancy and streamlines the process of provisioning multiple tunnels.

For more information, see ["Network Optimization Overview" on page 1294](#) and ["About the Network Optimization Intent Page" on page 1298](#).

Tasks You can Perform

- View details of a tunnel profile—You can select and view the details of a tunnel profile. You can either click **More > Detail** or hover over the tunnel profile and click the **Detailed View** icon. The Tunnel Profile Details pane appears on the right side of the page displaying the details.

Click the **Close (x)** icon to close the pane.

- Create a tunnel profile— Click **+** to create a tunnel profile. For more information, see ["Add a Tunnel Profile" on page 1305](#).
- Edit a tunnel profile—To edit a tunnel profile, select the tunnel profile and click the pencil icon. After you edit, you can choose to save the tunnel profile as a draft or publish the tunnel profile.



NOTE: You cannot edit a tunnel profile if it is linked to a path intent.

- Delete a tunnel profile—To delete a tunnel profile, select a tunnel profile and click the delete icon. A message confirms that Routing Director has successfully deleted the tunnel profile.



NOTE: You cannot delete a tunnel profile if it is linked to a path intent.

To perform common actions such as filtering using advanced filter criteria and resetting preferences, see ["Network Table Overview" on page 515](#).

Field Descriptions

[Table 318 on page 1305](#) lists the fields on the Tunnel Profiles tab.

Table 318: Fields on the Tunnel Profiles Tab

Column	Description
Name	Displays the name of the tunnel profile.
Description	Displays the description that you have entered at the time of tunnel profile creation.
Status	Displays the status (Draft or Published) of the tunnel profile. Only published tunnel profile can be attached to a path intent.
Linked Intents	Displays the number of path intents the tunnel profile is attached to.

Add a Tunnel Profile

Create tunnel profiles to define some of tunnel properties like signaling protocol, routing method, and so on. For more information, see ["Network Optimization Overview" on page 1294](#) and ["About the Tunnel Profiles Tab" on page 1304](#).

To create a tunnel profile:

1. Navigate to **Network Optimization > Network > Network Optimization Intent**.

The Topology page is displayed with the topology map at the center and the network information table at the bottom of the page.

2. Navigate to the Tunnel Profiles tab.

3. Click the plus icon (+).

The Create Tunnel Profile page appears.

4. Complete the configuration according to the guidelines in [Table 319 on page 1306](#).



NOTE: Fields marked with an asterisk (*) are mandatory.

5. Do one of the following:

- Click **Save Draft**, to save the tunnel profile. Save the tunnel profile if you intend to modify it later.
- Click **Publish**, to publish the tunnel profile. Publish the tunnel profile if you want to attach this tunnel profile to a path intent.



NOTE: You can attach the tunnel profile to a path intent only if you publish the tunnel profile. If a tunnel profile is attached to a path intent, then the tunnel profile cannot be edited or deleted.

A message indicating that the tunnel profile is created appears. You are returned to the Tunnel Profiles tab where you can view the newly added tunnel profile.

Table 319: Fields on the Create Tunnel Profile Page

Column	Description
Name*	Specify the name of the tunnel profile. You cannot use space, comma, or semicolon.
Description	Enter a description for the tunnel profile.

Table 319: Fields on the Create Tunnel Profile Page (*Continued*)

Column	Description
Provisioning Method	From the list, select one of the following methods to be used to provision the tunnel: • PCEP (Path Computation Element Protocol)—The path computation element (PCE) initiates the tunnel and the associated configuration statements do not appear in the router configuration file. Upon provisioning, this tunnel is added as a PCE-initiated tunnel. This is the default provisioning method. • NETCONF—The tunnel is statically provisioned and the associated configuration statements appear in the router configuration file. Upon provisioning, this tunnel is added as a device-controlled tunnel. NOTE: • For Cisco IOS-XR routers, NETCONF-based tunnel provisioning has the same capabilities as PCEP-based tunnel provisioning. • When provisioning tunnels by using NETCONF one at a time, the provisioning order might be sent before the response to a previous provisioning order is received. The second order might not have the correct bandwidth allocation information and the PCE might not be able to provide ECMP. We recommend provisioning multiple tunnels through NETCONF in one operation (bulk provisioning) in order to avoid this issue. • Device-controlled tunnels are not optimized during the period specified in the optimization profile.
Signaling protocol	Select the signaling protocol (RSVP or SR) to provision the LSP.

Table 319: Fields on the Create Tunnel Profile Page (*Continued*)

Column	Description
Routing Method	Select default, delay, IS-IS or OSPF if you want the PCE to compute and provision the path as a strict explicit route. Select Route by device, if you want PCE to learn or create a PCC-controlled tunnel. For this method, the PCE does not compute and provision a path.
Properties	Click the Properties button to configure additional attributes such as planned bandwidth, priority, color community for a tunnel. The Tunnel Profile Properties page appears.
Tunnel Profile Properties	This page appears when you click the Properties button. Select the properties that you want to configure for a tunnel. The properties that you select appear on the Create Tunnel Profile page. You can specify the values for the properties on the Create Tunnel Profile page.
Planned Bandwidth	Specify the planned bandwidth for the tunnel. By default, the unit is bps.
Setup Priority	Specify the setup priority for the tunnel traffic. Priority levels range from 0 (highest priority) through 7 (lowest priority). The default is 7, which is the standard MPLS tunnel definition in Junos OS. The PCE uses the setup priority to determine whether a new tunnel can be established, by preempting an existing tunnel. The existing tunnel can be preempted if the setup priority of the new tunnel is higher than that of the existing tunnel and the preemption releases enough bandwidth for the new tunnel.

Table 319: Fields on the Create Tunnel Profile Page (*Continued*)

Column	Description
Hold Priority	Specify the hold priority for the tunnel traffic. Priority levels range from 0 (highest priority) through 7 (lowest priority). The default is 7, which is the standard MPLS tunnel definition in Junos OS. The PCE uses the hold priority to determine whether the tunnel can be preempted or not. If the hold priority for a tunnel is higher, it is unlikely for the tunnel to be preempted.
Color Community	Assign a color to a segment routing tunnel. A segment routing (SR) policy uses the color community value to steer traffic flows on to a tunnel
FlexAlgoID	Select the flex algo ID that you want to associate with the LSP. This field lists all the flex algo IDs in your network. When the LSP is provisioned, the LSP adheres to the pre-defined FAD constraints such as administrative group restrictions, metric, optimizing traffic paths, and so on. You can select the flex algo ID only for SR (segment routing) and SRv6 LSPs.

RELATED DOCUMENTATION

Path Intent Workflow | 1296

About the Optimization Profiles

IN THIS SECTION

- [Tasks You can Perform | 1310](#)
- [Field Descriptions | 1311](#)

To access the Optimization Profiles tab, click **Network Optimization > Network > Network Optimization Intent**.

On this tab, you can create optimization profiles, view the status (Published or Draft) and the number of intents linked to the optimization profile.

Optimization profiles define when the LSP must be recomputed, which can be later used to create path intents. For more information, see ["Network Optimization Overview" on page 1294](#) and ["About the Network Optimization Intent Page" on page 1298](#).

Tasks You can Perform

- View details of an optimization profile—You can select and view the details of an optimization profile. You can either click **More > Detail** or hover over the optimization profile and click the **Detailed View** icon. The Optimization Profile Details pane appears on the right side of the page displaying the details.

Click the **Close (x)** icon to close the pane.

- Create an optimization profile— Click **+** to create an optimization profile. For more information, see ["Add an Optimization Profile" on page 1311](#).
- Edit an optimization profile—To edit an optimization profile, select the optimization profile and click the pencil icon. After you edit, you can choose to save the optimization profile as a draft or publish the optimization profile.



NOTE: You cannot edit an optimization profile if it is linked to a path intent.

- Delete an optimization profile—To delete an optimization profile, select an optimization profile and click the delete icon. A message confirms that Routing Director has successfully deleted the optimization profile.



NOTE: You cannot delete an optimization profile if it is linked to a path intent.

To perform common actions such as filtering using advanced filter criteria and resetting preferences, see ["Network Table Overview" on page 515](#).

Field Descriptions

[Table 320 on page 1311](#) lists the fields on the Optimization Profile tab.

Table 320: Fields on the Optimization Profile Tab

Column	Description
Name	Displays the name of the optimization profile.
Description	Displays the description that you have entered at the time of optimization profile creation.
Status	Displays the status (Draft or Published) of the optimization profile. Only published optimization profile can be attached to a path intent.
Linked Intents	Displays the number of path intents the optimization profile is attached to.

Add an Optimization Profile

Create optimization profiles to define when the LSP must be recomputed. For more information, see ["Network Optimization Overview" on page 1294](#) and ["About the Optimization Profiles" on page 1310](#).

To create an optimization profile:

1. Navigate to **Network Optimization > Network > Network Optimization Intent**.

The Topology page is displayed with the topology map at the center and the network information table at the bottom of the page.

2. Navigate to the Optimization Profiles tab

- 3. Click the plus icon (+).

The Create Optimization Page page appears.

- 4. Complete the configuration according to the guidelines in [Table 320 on page 1311](#).



NOTE: Fields marked with an asterisk (*) are mandatory.

- 5. Do one of the following:

- Click **Save Draft**, to save the optimization profile. Save the optimization profile if you intend to modify it later.
- Click **Publish**, to publish the optimization profile. Publish the optimization profile if you want to attach this optimization profile to a path intent.



NOTE: You can attach the optimization profile to a path intent only if you publish the optimization profile. If an optimization profile is attached to a path intent, then the optimization profile cannot be edited or deleted.

A message indicating that the optimization profile is created appears. You are returned to the Optimization Profiles tab where you can view the newly added optimization profile.

Table 321: Fields on the Create Optimization Profile Page

Column	Description
Name*	Specify the name of the optimization profile. You cannot use space, comma, or semicolon.
Description	Enter a description for the optimization profile.
Time Interval*	Specify the time duration (minutes). A new optimized path is created within the specified time duration.

RELATED DOCUMENTATION

[Path Intent Workflow](#) | 1296

About the Endpoint Profiles

IN THIS SECTION

- [Tasks You can Perform | 1313](#)
- [Field Descriptions | 1314](#)

To access the Endpoint Profiles tab, click **Network Optimization > Network > Network Optimization Intent**.

On this tab, you can create endpoint profiles, view the status (Published or Draft) and the number of intents linked to the endpoint profile.

Endpoint profiles define the endpoints of the LSPs (tunnels), which can be later used to create a path intents. For more information, see ["Network Optimization Overview" on page 1294](#) and ["About the Network Optimization Intent Page" on page 1298](#).

Tasks You can Perform

- View details of an endpoint profile—You can select and view the details of an endpoint profile. You can either click **More > Detail** or hover over the endpoint profile and click the **Detailed View** icon. The Endpoint Profile Details pane appears on the right side of the page displaying the details.

Click the **Close (x)** icon to close the pane.
- Create an endpoint profile— Click **+** to create an endpoint profile. For more information, see ["Add an Endpoint Profile" on page 1314](#).
- Edit an endpoint profile—To edit an endpoint profile, select the endpoint profile and click the pencil icon. After you edit, you can choose to save the endpoint profile as a draft or publish the endpoint profile.



NOTE: You cannot edit an endpoint profile if it is linked to a path intent.

- Delete an endpoint profile—To delete an endpoint profile, select an endpoint profile and click the delete icon. A message confirms that Routing Director has successfully deleted the endpoint profile.



NOTE: You cannot delete an endpoint profile if it is linked to a path intent.

To perform common actions such as filtering using advanced filter criteria and resetting preferences, see ["Network Table Overview" on page 515](#).

Field Descriptions

[Table 322 on page 1314](#) lists the fields on the Endpoint Profiles tab.

Table 322: Fields on the Endpoint Profiles Tab

Column	Description
Name	Displays the name of the optimization profile.
Description	Displays the description that you have entered at the time of endpoint profile creation.
Status	Displays the status (Draft or Published) of the endpoint profile. Only published optimization profile can be attached to a path intent.
Linked Intents	Displays the number of path intents the endpoint profile is attached to.

Add an Endpoint Profile

Create endpoint profiles to define the endpoints for the LSP. For more information, see ["Network Optimization Overview" on page 1294](#) and ["About the Endpoint Profiles" on page 1313](#).

To create an endpoint profile:

1. Navigate to **Network Optimization > Network > Network Optimization Intent**.

The Topology page is displayed with the topology map at the center and the network information table at the bottom of the page.

2. Navigate to the Endpoint Profiles tab.

3. Click the plus icon (+).

The Create Endpoint Page page appears.

4. Complete the configuration according to the guidelines in [Table 323 on page 1315](#).



NOTE: Fields marked with an asterisk (*) are mandatory.

5. Do one of the following:

- Click **Save Draft**, to save the endpoint profile. Save the endpoint profile if you intend to modify it later.
- Click **Publish**, to publish the endpoint profile. Publish the endpoint profile if you want to attach this endpoint profile to a path intent.



NOTE: You can attach the endpoint profile to a path intent only if you publish the endpoint profile. If an endpoint profile is attached to a path intent, then the endpoint profile cannot be edited or deleted.

A message indicating that the endpoint profile is created appears. You are returned to the Endpoint Profiles tab where you can view the newly added endpoint profile.

Table 323: Fields on the Create Endpoint Profile Page

Column	Description
Name*	Specify the name of the endpoint profile. You cannot use space, comma, or semicolon.
Description	Enter a description for the endpoint profile.
Add Nodes from the Map	Click here to select the source and destination endpoints for the tunnels directly from the topology map.
Filter A Nodes	Filter the source nodes in the network using logical operators such as AND, OR, and NOT and conditions such as =, !=, <, <=, >, >=, IN, NOT IN, INCLUDES, and NOT INCLUDES.

Table 323: Fields on the Create Endpoint Profile Page *(Continued)*

Column	Description
Filter Z Nodes	Filter the destination nodes in the network using logical operators such as AND, OR, and NOT and conditions such as =, !=, <, <=, >, >=, IN, NOT IN, INCLUDES, and NOT INCLUDES.
Node A Selection	Displays the list of source endpoints that you have selected.
Node Z selection	Displays the list of destination endpoints that you have selected.
Connection	Select the type of connection between the endpoints. The connection type can be one-to-one, one-to-all, or full mesh. The one-to-one and one-to-all can be both unidirectional or bidirectional.

RELATED DOCUMENTATION

[Path Intent Workflow](#) | 1296

About the Device Tab (Network Optimization)

IN THIS SECTION

- [Tasks You Can Perform](#) | 1317

To access the Device tab, click **Network Optimization > Network > Network Optimization Intent**.

On this tab, you can view information (such as the hostname, IP address, and type) about the devices in the network. [Table 125 on page 518](#) provides information on the fields in the Device tab.

You can also perform various actions on the devices from this tab. To perform common actions such as filtering using advanced filter criteria, downloading device details, and resetting preferences, see ["Network Table Overview" on page 515](#).

Tasks You Can Perform

- From the **More** list, you can perform the following tasks:



NOTE: You can also right-click a device row in the table to view the same options.

- View details about a device—To view details (such as the properties configured for a device and its interfaces):
 - Hover over a device row in the table and click the **Details** icon. On the **Device *Device-Name*** page that appears, you can view the Details tab and the Interfaces tab.
 - Select the device on the table and click **More > Show Detail**. On the **Device *Device-Name*** page that appears, you can view the Details tab and the Interfaces tab.

On the Details tab, you can view general device properties such as, model, device type, MAC address, and the serial number. You can also view site properties such as site name, address, country, and time zone where the device is deployed. On the Interfaces tab, you can view all the interfaces of the device, the administrative and operational statuses, the speed of the physical interfaces (in Mbps), and the duplex mode. The **Device *Device-Name*** page is moveable, so you can pin it anywhere on the screen.



NOTE: Routing Director populates the speed column for physical interfaces and for interfaces with Operational Status Up.

- Filter a device in the topology map—To display only the selected device in the topology map, select **Filter Selected device**. The Devices & Links page reloads to display all devices and links on the topology map. To undo the filter, right click a blank space on the topology map and click **Reload Network**.
- Zoom in on a device—To zoom in on a device in the topology map, **select** the device and click **Zoom to Selected device**. The topology map is enlarged to zoom in on the selected device.
- Assign a device to a site—You can select devices and add them to a site. See ["About the Troubleshoot Devices Page" on page 454](#) for steps to assign devices to sites.

- Back up device configuration—You can back up device configuration. In the event of a device failure, you can use the backed up configuration to restore the device configuration. See ["About the Troubleshoot Devices Page" on page 454](#) for the steps to back up a device configuration.
- Reboot—You can reboot the device after making configuration changes or commit changes in a planned maintenance event. See ["About the Troubleshoot Devices Page" on page 454](#) for the steps to reboot the device.
- Upgrade—You can upgrade the current software image on the device to the latest version. See ["About the Troubleshoot Devices Page" on page 454](#) for the steps to upgrade the software image.

Table 324: Fields on the Device Table

Fields	Description
Hostname	Lists the physical and virtual hosts (devices) in your network.
Severity	Displays the highest severity of the events on the device. When you hover over the severity, a pop-up displays the total number of events for all severity levels. The severity levels are: • Urgent Action Needed or critical • Action Needed or major • Being Monitored or minor If there are no events, you can see a green ✓ icon, indicating a healthy device status.
IP Address	Displays the management IPv4 address of the physical and virtual devices in your network.
Site	Displays the geographical site to which you added the device.
Model	Displays the model of the device. For example, ACX-7100-32C.
Serial Number	Displays the serial number of the device.

Table 324: Fields on the Device Table *(Continued)*

Fields	Description
Latitude	Displays the latitude value of a device. Latitudes range from -90 to 90. Positive values of latitude are north of the equator and negative values (preceded with a minus sign) are south of the equator.
Longitude	Displays the longitude value of a device. Longitudes range from -180 to 180. Positive longitudes are east of the Prime Meridian and negative values (preceded with a minus sign) are west of the Prime Meridian.
OS Version	Version of the OS that is currently installed on the device.
PCEP Status	Display whether path computation element protocol (PCEP) is configuration on a device or not.
PCEP IP	Displays the IP address of the PCEP.
SID	Displays the segment identifier (SID) if you have configured segment routing.
SR	Display whether segment routing (SR) is configuration on a device or not.
SRLG	Displays information about shared risk link group (SRLG).
Live	Displays the status of the device.
FAD Definition	Displays the flexible algorithms definition (FAD) ID. Hover over the FAD ID link to view the flex algo configuration.
FAD Participation	Displays the list of FAD IDs that the device is part of.

RELATED DOCUMENTATION

[Network Optimization Overview | 1294](#)

[About the Network Optimization Intent Page | 1298](#)

About the Link Tab (Network Optimization)

IN THIS SECTION

- [Tasks You Can Perform | 1320](#)

To access the Link tab, click **Network Optimization > Network > Network Optimization Intent**.

On this tab, you can view information about the links in the network. [Table 126 on page 522](#) provides information on the fields in the Link tab.

To perform common actions such as filtering using advanced filter criteria, downloading link details, and resetting preferences, see "[Network Table Overview](#)" on page 515.

Tasks You Can Perform

- From the **More** list, you can perform the following tasks:



NOTE: You can also right-click a link to view the same options.

- View details about a link—To view details such as the site name, country, site id, organization, id, and time zone:
 - Hover over a link row in the table and click the **Details** icon. The **Link - Device A to Device Z** appears. On the **Link - Device A to Device Z** page that appears, you can view the Interface Stats tab and the Details tab.
 - Select the link on the table and click **More > Show Detail**. On the **Link - Device A to Device Z** page that appears, you can view the Interface Stats tab and the Details tab.

On the Interface Stats tab, you can view the device name, interface name, and management IP address of device A and device Z. On the Details tab, you can view the node id associated with the interfaces, interface IPv4 address at the two nodes, interface name, interface bandwidth, live

status, and the operational status of the interface. The **Link - Device A to Device Z** page is movable so, you can pin it anywhere on the screen.

- Filter links on the topology map—To display only the selected link in the topology map, select **Filter Selected Link**. To undo the filter, right click a blank space on the topology map and click **Reload Network**.
- Zoom in on a link—To zoom in on a link in the topology map, select the link and click **Zoom to Selected Link**. The topology map is enlarged to zoom in on the selected link.
- Configure table settings—Select **Table Settings** to specify the link utilization format (**Decimal** or **Percentage**) for the link utilization parameters in the specified format.
- Download link-related information—Click **Download** to download detailed information about all the existing links in the topology in CSV format.
- View historical data—Routing Director enables you can view the historical data. Historical data provides you valuable insights into past trends and patterns.



NOTE: For Routing Director to collect historical data, you need to do the following:

1. You need to enable the collection of statistics.
 - a. To collect historical data for interface traffic, you need to enable traffic engineering while creating a device profile. For more information, see ["Add a Device Profile" on page 208](#).
 - b. To collect historical data for interface delay, you need to enable link delay measurements. For more information, see <https://www.juniper.net/documentation/us/en/software/junos/is-is/topics/topic-map/enable-link-delay-advertise-in-is-is.html>
2. Enable dynamic topology so that you can view real-time changes in your network. For more information, see ["Dynamic Topology Workflow" on page 555](#).

You can view the historical data for the following:

- Interface Traffic—View a graphical representation of the link traffic (in bps) for interface A and Z based on the time range selected from the **Period** list. A **Link Traffic** pop-up appears, which you can pin anywhere on the screen using the pin icon. From the **Period** list, you can select an option to view the traffic data for Previous 3 Hours, 1 day, 1 Week, or Custom time ranges.

You can choose to auto refresh the data or manually refresh the data.

- Interface delay—View a graphical representation of the link delay (in milliseconds) for interface A and Z based on the time range selected from the **Period** list. A **Link Delay** pop-up appears, which

you can pin anywhere on the screen using the pin icon. From the **Period** list, you can select an option to view the delay data for Previous 3 Hours, 1 day, 1 Week, or Custom time ranges.

You can choose to auto refresh the data or manually refresh the data.

Table 325: Fields on the Link Table

Fields	Description
Device A	Displays the device at which traffic enters.
Device Z	Displays the device at which traffic exits.
Status	Displays the status of the link.
IP A	Displays the IPv4 address of the interface from which device A sends traffic. Routing Director displays the IPv4 address of IP A based on device A's active configuration. If the IPv4 addresses of IP A and IP Z are in the same subnet, Routing Director forms an interface between devices A and Z.
IP Z	Displays the IPv4 address of device Z interface that receives traffic. Routing Director displays the IPv4 address of IP Z based on device Z's active configuration. If the IPv4 addresses of IP A and IP Z are in the same subnet, Routing Director forms an interface between devices A and Z.
Interface A	Displays the interface name of device A.
Interface Z	Displays the interface name of device Z.
Bandwidth A	Displays the bandwidth (in bps) of the ingress device.
Bandwidth Z	Displays the bandwidth (in bps) of the egress device.
TE Metric A	Displays the traffic engineering-related metrics for the device A.

Table 325: Fields on the Link Table *(Continued)*

Fields	Description
TE Metric Z	Displays the traffic engineering-related metrics for the device B.
Admin Group AZ	Displays the admin group selected for the traffic to traverse between device A and device Z.
Admin Group ZA	Displays the admin group selected for the traffic to traverse between device Z and device A.
Interface Util AZ	Displays the interface utilization between device A and device Z.
Interface Util ZA	Displays the interface utilization between device Z and device A.
RSVP BW AZ	Displays the RSVP bandwidth between device A and device Z.
RSVP BW ZA	Displays the RSVP bandwidth between device Z and device A.
RSVP Util AZ	Displays the RSVP utilization between device A and device Z.
RSVP Util ZA	Displays the RSVP utilization between device Z and device A.
Measured Delay AZ	Displays the delay metrics between device A and device Z.
Measured Delay ZA	Displays the delay metrics between device Z and device A.
SRLG	Displays information about shared risk link group (SRLG).
SID A	Displays the SID label for the ingress device.

Table 325: Fields on the Link Table *(Continued)*

Fields	Description
SID Z	Displays the SID label for the egress device.
ISIS1 AZ	Displays the IS-IS area level (L1) metric between device A and device Z.
ISIS1 ZA	Displays the IS-IS area level (L1) metric between device Z and device A.
ISIS2 AZ	Displays the IS-IS area level (L2) metric between device A and device Z.
ISIS2 ZA	Displays the IS-IS area level (L2) metric between device Z and device A.
OSPF AZ	Displays the OSPF metric between device A and device Z.
OSPF ZA	Displays the OSPF metric between device Z and device A.

RELATED DOCUMENTATION

[Network Optimization Overview | 1294](#)

[About the Network Optimization Intent Page | 1298](#)

About the Tunnels Tab (Network Optimization)

IN THIS SECTION

- [Tasks You Can Perform | 1325](#)

To access the Tunnels tab, click **Network Optimization > Network > Network Optimization Intent**.

On the Tunnels tab, you can view the information about all label-switched paths (LSPs) or segment routing policies, also known as tunnels, in your network. Routing Director uses Path Computation Element Protocol (PCEP) sessions to discover the tunnels that run in your network.



NOTE: For Routing Director to discover tunnels in your network:

1. A PCEP session must be established between Routing Director and the devices.

To establish a PCEP session, you must specify a VIP address either at the time of installing Routing Director or post-installation of the cluster. For more information, see ["Configure a PCE Server" on page 556](#).

To establish PCEP on a router, see ["Configuring PCEP on a PE Router" on page 557](#).

To configure BGP-LS topology acquisition on Routing Director, see ["Configuring Topology Acquisition Using BGP-LS" on page 559](#).

2. You must have adopted or onboarded the device using Routing Director; otherwise, the tunnel-related information is not displayed on the Tunnels tab.

[Table 326 on page 1327](#) provides information on the fields in the Tunnels tab.

To perform common actions such as filtering using advanced filter criteria and resetting preferences, see ["Network Table Overview" on page 515](#).

Tasks You Can Perform



NOTE: You cannot provision a tunnel from the Tunnels tab of the Network Optimization Intent page (**Network Optimization > Network**). You must provision the tunnel from the Tunnels tab of the Topology page (**Observability > Network**). For more information see, ["Add a Tunnel" on page 535](#).

- Hide unrelated nodes of a tunnel—To hide unrelated nodes of a tunnel, select one or more tunnels and enable the **Hide unrelated** toggle button. The topology map limits the display to only the nodes that are related to the selected tunnel.
- Provision a tunnel—Click **Provisioning** and select **Tunnel** to provision a tunnel. For more information, see ["Add a Tunnel" on page 535](#).



NOTE: Only Path Computation Element Protocol (PCEP) provision method is supported.

- View historical data—Routing Director enables you to view the historical data. Historical data provides you valuable insights into past trends and patterns.



NOTE: For Routing Director to collect historical data, you need to do the following:

1. You need to enable the collection of statistics.
 - a. To collect historical data for tunnel traffic, you need to enable traffic engineering while creating a device profile. For more information, see ["Add a Device Profile" on page 208](#).
 - b. To collect historical data for tunnel delay, you need to enable link delay measurements. For more information, see <https://www.juniper.net/documentation/us/en/software/junos/is-is/topics/topic-map/enable-link-delay-advertise-in-is-is.html>
2. Enable dynamic topology so that you can view real-time changes in your network. For more information, see ["Dynamic Topology Workflow" on page 555](#).

You can view the historical data for the following:

- Tunnel traffic—To view traffic and bandwidth for a tunnel in graphical form, select the tunnel and click **Tunnel Traffic**. In the Tunnel Traffic page that appears, select the period for which you want to view the data. You can view data for the previous 3 hours, the previous day, the previous week, or choose from a custom time range by specifying the start and end dates and the time.



NOTE: You can view the tunnel traffic if the device has Junos OS Release 22.X or later releases.

If you are using ACX Series device, then you can view tunnel traffic only if you have installed Junos OS Evolved Release 23.4R2-S2.1 or later releases.

- Tunnel delay—Tunnel delay is the sum of all link delays in the tunnel path. At any given time, the Path Computation Element (PCE) is aware of the paths of all tunnels in the network. Periodically, the PCE uses the reported link delays to compute the end-to-end tunnel delay.

To view tunnel delay in graphical form, select the tunnel and click **Delay**. In the Tunnel Delay page that appears, select the period for which you want to view the data. You can view data for the previous 3 hours, the previous day, the previous week, or choose from a custom time range by specifying the start and end dates and the time.

You can choose to refresh the data automatically or manually.

- View events for tunnels—To view historical events (such as actions performed on the tunnel and bandwidth changes) for a tunnel for a specific time range, select the tunnel and click **View > Event History**.

On the Events page, you can:

- View a graphical representation of bandwidth changes and traffic through an LSP to analyze trends over time. The bandwidth changes represent the changes in the signaled bandwidth.

Select the period for which you want to view the bandwidth-related data. You can view data for the previous 2 hours, 4 hours, 8 hours, 16 hours, 24 hours, 1 week, or choose from a custom time range by specifying the start and end dates and the time.

- View changes to label-switched path (LSP) attributes and events in the table. The Timestamp column helps you to identify the time at which a change has been made to a tunnel.
- View changes to label-switched path (LSP) on the Topology map. You can use the **More > Show Path Changes** option to view highlights on the Topology map for any changes to the LSP path, making it easier to track and analyze routing changes.
- Download comma-separated values (CSV) file that includes information on LSP event history to your local system by clicking the **More > Show Path Changes** option.
- Download—To download tunnels-related data to your local system, click **Download**. The data is downloaded to your local system as a comma-separated values (CSV) file.
- From the More list, you can perform the following tasks:
 - View details of a tunnel—To view the details of a tunnel, you can do one of the following:
 - Select a tunnel on the table and click **More > Show Details**. The **Tunnel - Tunnel-Name** page appears. You can view the To and From traffic-related information.
 - Hover over a tunnel row in the table and click the **Details** icon. The **Tunnel - Tunnel-Name** page appears. You can view the To and From traffic-related information.
 - Reload—To view the latest changes in the network map, click **Reload**.
- To perform common actions such as filtering using advanced filter criteria and resetting preferences, see "[Network Table Overview](#)" on page 515.

Table 326: Fields on the Tunnel Table

Fields	Description
Name	Displays the name of the tunnel.

Table 326: Fields on the Tunnel Table *(Continued)*

Fields	Description
Device A	Displays the device at which traffic enters.
Device Z	Displays the device at which traffic exits.
IP A	Displays the IPv4 address of the interface from which device A sends traffic. Routing Director displays the IPv4 address of IP A based on device A's active configuration. If the IPv4 addresses of IP A and IP Z are in the same subnet, Routing Director forms an interface between devices A and Z.
IP Z	Displays the IPv4 address of device Z interface that receives traffic. Routing Director displays the IPv4 address of IP Z based on device Z's active configuration. If the IPv4 addresses of IP A and IP Z are in the same subnet, Routing Director forms an interface between devices A and Z.
Vendor	Displays the vendor name of the device.
Bandwidth	Displays the traffic rate associated with the LSPs, fast reroutes, or multiclass LSPs.
Color	Displays the color associated with the segment routing (SR) policy. A segment routing (SR) policy uses the color community value to steer traffic flows on to a tunnel.
Admin Group Include All	Displays a list of all admin groups that are created for LSPs. Admin groups are created to implement a variety of policy-based LSP setups.

Table 326: Fields on the Tunnel Table *(Continued)*

Fields	Description
Admin Group Include Any	Displays a list of the admin groups that are created for LSPs. Admin groups are created to implement a variety of policy-based LSP setups.
Metric	Displays the metric value of an LSP. The LSP metric is used to indicate the ease or difficulty of sending traffic over a particular LSP. Lower LSP metric values (lower cost) increase the likelihood of an LSP being used. Conversely, high LSP metric values (higher cost) decrease the likelihood of an LSP being used.
Path Type	Displays the type of LSP path—primary, secondary, or standby.
Path Name	Displays the user-defined name of the LSP path.
Type	Displays the type of the LSP—RSVP or SR.
Op Status	Displays the operational status of the LSP—Failed, pending, unknown, active.
Setup	Displays the setup priority of the LSP. The setup priority ranges from 0 (highest) through 7 (lowest).
Hold	Displays the RSVP hold priority for the tunnel traffic. The hold priority ranges from 0 (highest) through 7 (lowest).

Table 326: Fields on the Tunnel Table *(Continued)*

Fields	Description
Record Route	Displays whether an LSP should actively record the path by sending the record route object (RRO). The RRO is used to record the path that the LSP traverses. It includes the IP address and router/node ID of each RSVP node in the path.
Last Update	Displays the last updated time of the LSP.
Controller Status	Displays the controller status—No path found, Reschedule in <i>X</i> minutes.
RRO Hop Count	Displays the hop count of the record route object (RRO). The RRO keeps track of the actual path the LSP is traversing.
Comment	Displays the comment that the user has added at the time of tunnel creation.
Admin Group Exclude	Displays the admin group that is excluded. The tunnel traverses links that do not include any of the admin groups listed here.
Planned Metrics	Displays the static tunnel metric. The PCE uses this metric to route the tunnel instead of allowing the router to choose a path
Signaling address	Displays the signaling address Egress Peer Engineering (EPE).
Path Intent	Displays the name of the path intent, if any.
Control Type	Displays the control type.

RELATED DOCUMENTATION

[Network Optimization Overview | 1294](#)

[About the Network Optimization Intent Page | 1298](#)

9

PART

Network Planner

- [Introduction | 1333](#)
 - [Offline Models | 1336](#)
 - [Offline Topology Map | 1343](#)
-

Introduction

IN THIS CHAPTER

- [Planner Overview | 1333](#)
- [Planner Workflow | 1334](#)

Planner Overview

IN THIS SECTION

- [Benefits of Planner | 1333](#)

The Planner use case is a planning and simulation tool that provides in-depth network views and reports on how the network is performing both during normal and failure scenarios, all without impacting your environment.

Planner can be used for offline visualization and detailed architectural planning of any production network. Planner enables you to forecast the impact that changes such as latency, additional traffic, shifts in traffic flows, and new capacity or services will have on your network.

Planner can generate a graphical view of a network, enabling the addition, removal, and reconfiguration of network elements. Using the network topology view, you can model and visualize dynamic, explicit routing paths, designed to operate within end-user defined constraints.

The effects of these changes and other traffic scenarios can be simulated without affecting the production network.

Benefits of Planner

- Plan for future growth by analyzing the current network capacity and future requirements for bandwidth. This analysis using Planner helps you to formulate plans for expanding your production

network to meet your business needs and simulate those plans before implementing the plan on the production network.

- Ensure network resiliency by studying the impact of node and link, and planning for mitigating any issues that may occur.
- Reduce network operation costs by formulating network designs for effective use of existing network resources.

RELATED DOCUMENTATION

| [Planner Workflow](#) | 1334

Planner Workflow

Use the Planning menu to view the offline network and its elements on the Topology map. In addition, you can analyze the impact of adding new demands or traffic before they are placed onto the real network and simulate failure scenarios.

The following workflow explains how to use Planner to simulate network failures and analyze response:

1. Get an offline version of your live network in the JSON or .gz format. In an offline network, you can make changes and forecast the impact without impacting the live network. To get an offline version of your network do one of the following:
 - a. Import your live network directly by clicking **More > Import from Live Network** on the Working Models page (**Planning > Networks > Offline Models**). For more information, see ["About the Working Models Page"](#) on page 1336.
 - b. Upload the JSON or .gz format file of the live network from your local system by clicking the **+** icon on the Working Models page. For more information see ["Create an Offline Network"](#) on page 1338.

The offline version of your live network is listed on the Working Models page.

2. On the Working Models page, select the imported or uploaded file (offline network) and click **More > Save Model** to save network-related data to a *runner* pod so that the latest changes are reflected in the offline network.
3. On the Working Models page, select the imported or uploaded file (offline network) and click the **Open** button. The topology map of your offline network is displayed.

4. On the offline topology map, make changes to the offline network by adding devices, links, tunnels, demands, and interfaces. Click **More > Save Model** option. For more information, see ["About the Offline Topology Page" on page 1343](#).
5. On the Working Models page, select the modified network file, and click the **Simulation** button. The Exhaustive Failure Simulation page is displayed,

Use the Exhaustive Failure Simulation page to create a failure scenario in your offline network and to see how your network reacts to the failure scenario. For information, see ["Exhaustive Failure Simulation" on page 1339](#).

A report is generated after you run the simulation.
6. On the Working Models page, click the **Reports** button to understand how your network reacted to the failure scenario. For more information, see ["View Reports" on page 1340](#).

RELATED DOCUMENTATION

| [Planner Overview](#) | **1333**

Offline Models

IN THIS CHAPTER

- [About the Working Models Page | 1336](#)
- [Create an Offline Network | 1338](#)
- [Exhaustive Failure Simulation | 1339](#)
- [View Reports | 1340](#)

About the Working Models Page

IN THIS SECTION

- [Tasks You Can Perform | 1336](#)
- [Field Descriptions | 1337](#)

To access the Working Models page, select **Planning > Networks > Offline Models**.

This page lists the offline versions of your live network (also called as network models) so that you can make changes to the network model and forecast the impact of the changes, without impacting the live network.

Use this page for offline visualization and detailed architectural planning of your production (live) network. You can create an offline network, generate a graphical view of the network, simulate failure scenarios, and view reports.

Tasks You Can Perform

- View the list of network models that have been uploaded or imported.

- Generate a topology map of a model—To generate the topology map of a model, select the model and click the **Open** button. You can view the topology map. For more information, see ["About the Offline Topology Page" on page 1343](#).
- Simulate error scenarios—Simulate failure scenarios in the network and to see how your network reacts, select the network model and click the **Simulation** button. For information, see ["Exhaustive Failure Simulation" on page 1339](#).
- View Reports—To understand how your network reacted to a failure scenario, click the **Reports** button. For more information, see ["View Reports" on page 1340](#).
- Create a model—To upload the JSON or .gz format file of the live network from your local system, click the + icon. For more information see ["Create an Offline Network" on page 1338](#).
- Save a network model—To save a network model and changes made to the network model on to the *runner* pod, click **More > Save Model**.
- Export a model—To export a model to your local system, click **More > Export**. The file is saved to your local system.
- Import from live network—To import a live network file directly to Routing Director, click **More > Import From Live Network**.
- You can also perform the following tasks on this page:
 - Sort, resize, or re-arrange columns in a table (grid).
 - Show or hide columns in the table or reset page preferences, using the vertical ellipsis menu.
 - Search by using keywords—Click the search icon (magnifying glass), enter the search term in the text box, and press Enter. The search results are displayed on the same page.
 - Filter the data displayed in the table—Click the filter icon (funnel) and select whether you want to show or hide advanced filters. You can then add or remove filter criteria, save criteria as a filter, apply or clear filters, and so on. The filtered results are displayed on the same page.
 - Set the number of rows (15, 30, 50, or 100) to be displayed in the table in the **Display** drop-down list provided at the bottom right-corner of the table. If there are multiple pages of rows, use the arrow heads to traverse through the pages or enter a page number to view the entries listed on that page.
 - For more information, see ["GUI Overview" on page 6](#).

Field Descriptions

[Table 327 on page 1338](#) lists the fields on the Working Models page.

Table 327: Fields on the Working Models Page

Field	Description
Topo ID	Displays a unique identification number for the network model.
Name	Displays the name of the model.
Description	Displays the description of the network model specified at the time of creating the network model.
Created by	Displays the name of the user who created the model.
Last Modified	Displays the date and time the network model was last modified.
Git Branch	Displays the Git branch number in which the network model is saved.

RELATED DOCUMENTATION

[Planner Overview](#) | 1333

[Planner Workflow](#) | 1334

Create an Offline Network

You must have an offline view of your live network in JSON or .gz format. In an offline network, you can make changes and forecast the impact without impacting the live network.

Use the following procedure to create an offline network:

1. Navigate to **Planning > Networks > Offline Models**.
The Working Models page is displayed.
2. Create the + icon to upload an offline version of your live network.
The Create Network Page is displayed.

3. Click the **Browse** button in the **Upload JSON/GZ File** field to upload the live network file from your local system in JSON and .gz format. The **Name** and **Description** fields are automatically updated after you upload the file.
4. Click the **Save** button to save the offline version of your live network.

The offline version of your live network (also called as a network model) is listed on the **Working Models** page. You can use this network model to view the offline topology map and make necessary changes to the network model.

Alternatively, you can import your live network by clicking **More > Import from Live Network** on the **Working Models** page. For more information, see ["About the Working Models Page" on page 1336](#).

RELATED DOCUMENTATION

[About the Working Models Page | 1336](#)

[Planner Workflow | 1334](#)

Exhaustive Failure Simulation

Use Routing Director to create a failure scenario in your network model. A report is generated after you run the simulation. Check the reports to understand the impact of various failures to the network.

To create a failure scenario:

1. Navigate to **Planning > Networks > Offline Models**.
The **Working Models** page is displayed.
2. From the list of network models (offline version of a live network), select a network model on which you want to simulate a failure scenario.
3. Click the **Simulation** button to define a failure scenario that you want to apply to simulate in your network model. The **Exhaustive Failure Simulation** page is displayed.



NOTE: Before you run a simulation, ensure that you save the changes made to the network model by clicking **More > Save Model**.

4. Complete the configuration according to the guidelines in [Table 328 on page 1340](#).



NOTE: Fields marked with an asterisk (*) are mandatory.

5. Click **Simulate** to create a failure scenario.

A message indicating that the simulation is complete is displayed. A report is generated after the simulation is complete. On the Reports page (**Working Models** > select the model > click **Reports**), check the reports to understand how your network reacted to the failure scenario. For more information, see ["View Reports" on page 1340](#).

Table 328: Fields on the Exhaustive Failure Simulation Page

Field	Description
Layer*	Select the layer for which you want to create the failure scenario. You can choose one of the following: • Demand (Layer 3) • Tunnel (Layer 2)
Description	Describe the failure scenario that you are creating.
Exhaustive Failure*	Select the type of failure (single) you want to simulate. Currently, you can simulate and check how the network is affected when one of the links or nodes in the network fails at a time. For example, if a network has 10 links, at any point in time, Routing Director fails only one link and checks how the network reacts to the failure scenario. The same is applicable for the node failure.
Failure Type*	Select the failure type. You can simulate failure scenarios for links and nodes.

RELATED DOCUMENTATION

[Planner Overview](#) | 1333

[Planner Workflow](#) | 1334

[View Reports](#) | 1340

View Reports

Planner offers extensive reporting capabilities for network planning and simulation. These reports help quickly assess and summarize network-related data and can be generated in the CSV (comma-separated) or text format.

After you run simulation on your network model, Routing Director generates a report. These reports are listed on the Reports page. (**Planning > Networks > Offline Models > *Model-Name* > Reports**).

On the **Reports** page, you can do the following:

- View the list of reports generated for each simulation in the Simulation table. The Simulation table displays the following information (columns):
 - Report—Every time you run a simulation, a list of reports is generated. This column groups the reports simulation-wise. Click on each report to view detailed information about that report.
 - Description—Describes the report.
 - Created—Displays the date and time at which the report is generated.
 - Filename—Displays the filename of the report. Click the link to download the report.
 - Size—Size of the report.
- View the details of each report— To view the details of each report, click the **Exhaustive** down-arrow in the Simulation section. The following reports are generated for tunnels and demands:
 - Traffic on Links—Includes information on each link's peak utilization caused due to one or more failed elements. Click the link to view the report.
 - Oversubscribed Links—Includes information on the links that have more than 100% utilization during exhaustive failure simulation. Click the link to view the report.
 - Failed Tunnel Path —Includes information on tunnels that were unable to reroute, and the events that prevented tunnel reroute during exhaustive failure simulation. Click the link to view the report. This report is generated only for tunnels .
 - Failed Path —Includes information on failed paths, and the events that prevented reroute during exhaustive failure simulation. Click the link to view the report. This report is generated only for demands.
 - Simulation Summary—Includes information on the summary of the count, bandwidth, and the number of hops of the tunnels that were failed or were impacted during exhaustion simulation. Click the link to view the report.
- Download reports—You can download reports in the following ways:
 - On the Simulation table, click the link in the Filename column. The report is downloaded to your local system.

- On the Simulation table, click the **Exhaustive** down-arrow > *Report-Name* (link). You are directed to the page that displays information on the specific report. Click the **Download** option. The report is downloaded to your local system.

RELATED DOCUMENTATION

[Planner Overview | 1333](#)

[Planner Workflow | 1334](#)

[Exhaustive Failure Simulation | 1339](#)

Offline Topology Map

IN THIS CHAPTER

- [About the Offline Topology Page | 1343](#)
- [About the Device Tab \(Planner\) | 1345](#)
- [Create a Node \(Planner\) | 1347](#)
- [About the Link Tab \(Planner\) | 1348](#)
- [Create a Link \(Planner\) | 1352](#)
- [About the Tunnel Tab \(Planner\) | 1355](#)
- [Create a Tunnel \(Planner\) | 1357](#)
- [About the Demands Tab \(Planner\) | 1360](#)
- [Create a Demand \(Planner\) | 1362](#)
- [About the Interface Tab \(Planner\) | 1363](#)
- [Create an Interface \(Planner\) | 1365](#)

About the Offline Topology Page

To access the Offline Topology page:

1. Select **Planning > Networks > Offline Models**.
2. Select a network model and click **Open** to view the offline topology map .

On this page, you can view the topology map of your network model, modify the network model by adding network elements (devices, links, demands, interfaces, and tunnels), save the changes that you made to the network model, and run simulation for the failure scenarios that you have created.

The offline topology map page can be divided into three main components:

- **Interactive Topology Map**—Enables you to view devices, links, tunnels, and demands across sites in your network and customize the map display. Manage the map view, set device positions on the map, customize device, link labels, and font size using the right-click menu.

For more information, see ["View Network Topology Details" on page 504](#).

- **Topology Menu Bar**—A vertical bar at the top-right corner of the Network Optimization Intent page, which consists of the following icons:
 - Reset icon—Center the topology map so that it zooms to fit the screen.
 - Minus icon—Zoom out (reduce) the topology map.
 - Plus icon—Zoom in (enlarge) the topology map.
 - Switch to Cluster View/Node View icon—Switch to Cluster View from the default Node View and switch back to node view from cluster view.
 - Settings icon—Configure display properties of the topology map.

For more information, see ["View Network Topology Details" on page 504](#)

- **Network Table**—Displays detailed information about tunnels, demands, and network elements (devices, links, and interfaces).

[Table 329 on page 1344](#) lists the tabs on the offline topology map page.

Click the collapsible arrow icon at the bottom-left of the topology map to display or hide the network table.

Table 329: Tabs on the Offline Topology Map Page

Tab	Description
Device	See "About the Device Tab (Planner)" on page 1345 .
Link	See "About the Link Tab (Planner)" on page 1348 .
Tunnels	See "About the Tunnel Tab (Planner)" on page 1355 .
Demands	See "About the Demands Tab (Planner)" on page 1360 .
Interfaces	See "About the Interface Tab (Planner)" on page 1363 .

RELATED DOCUMENTATION

[Planner Overview | 1333](#)

[Planner Workflow | 1334](#)

About the Device Tab (Planner)

IN THIS SECTION

- [Tasks You Can Perform | 1345](#)

On the Device tab of the network table (**Planning > Networks > Offline Models > *Model-Name* > Open**), view information (such as the hostname, IP address, and type) about the devices in the offline network. [Table 330 on page 1346](#) provides information on the fields in the Device tab.

To perform common actions such as filtering using advanced filter criteria, downloading device details, and resetting preferences on the network table, see "[Network Table Overview](#)" on page 515.

Tasks You Can Perform

You can perform the following tasks on the Device tab:

- Download the device details by clicking the **Download** option. A CSV file is downloaded to your local system, which includes information on the devices in your network and its details.
- From the **More** list, you can perform the following tasks:



NOTE: You can also right-click a device row in the table to view the same options.

- View details about a device—To view details (such as the properties configured for a device):
 - Hover over a device row in the table and click the **Details** icon.
 - Select the device on the table and click **More > Show Detail**.

On the **Device *Device-Name*** page that appears, you can view the Details tab.

The Details tab includes information on device properties such as, model, device type, IP address, and the AS number.

The **Device** *Device-Name* page is moveable, so you can pin it anywhere on the screen.

- Add devices (node) to the network model—Click + to add new devices. For more information see ["Create a Node \(Planner\)" on page 1347](#).

Table 330: Fields on the Device Table

Fields	Description
Name	Name of the device.
Hostname	Lists the physical and virtual devices in your network.
IPv4 Address	Displays the management IPv4 address of the physical and virtual devices in your network.
Type	Displays the model of the device. For example, ACX-7100-32C.
Latitude	Displays the latitude of a device. Latitudes range from -90 to 90. Positive latitude indicate north of the equator and negative values (preceded with a minus sign) indicate south of the equator.
Longitude	Displays the longitude of a device. Longitudes range from -180 to 180. Positive longitudes indicate east of the Prime Meridian and negative values (preceded with a minus sign) indicate west of the Prime Meridian.
Description	Displays the description that you have entered at the time of node creation.
AS	Displays the AS number to which the node belongs to.
ISIS Area	Displays the ISIS area of the node.
ISIS System ID	Displays the ISIS system ID of the node.

RELATED DOCUMENTATION

| [Planner Workflow](#) | [1334](#)

Create a Node (Planner)

On the Node tab, add to device to your network model.

To add a device:

1. Navigate to **Planning > Networks > Offline Models > *Model-Name* > Open.**

The Topology page is displayed with the topology map at the center and the network information table at the bottom of the page.

2. Navigate to the Device tab.
3. Click the icon (+).

The Create Node page appears.

4. Complete the configuration according to the guidelines in [Table 331 on page 1347](#).



NOTE: Fields marked with an asterisk (*) are mandatory.

5. Click **Submit** to add a device.

A message indicating that the device is added appears. You are returned to the Device tab where you can view the newly added device.

Table 331: Fields on the Create Device Page

Field	Description
Name*	Specify the name of the device. Use alphanumeric characters, underscore, hyphen, period, or colon. Space, comma, semicolon not allowed.
Hostname	Specify the host name of the device.
Hardware	Select the device vendor.

Table 331: Fields on the Create Device Page (*Continued*)

Field	Description
Latitude	Specify the latitude of your device's location to visualize its placement on the topology map.
Longitude	Specify the longitude of your device's location to visualize its placement on the topology map
Properties	
Description	Enter a description for the node that you are adding.
IPv4	Specify the IPv4 address of the device.

RELATED DOCUMENTATION

[About the Offline Topology Page | 1343](#)

[Planner Workflow | 1334](#)

About the Link Tab (Planner)

IN THIS SECTION

- [Tasks You Can Perform | 1349](#)

On the Link tab (**Planning > Networks > Offline Models > *Model-Name* > Open**), view information about the links in the network. You can also perform various actions on the links from this tab. To perform common actions such as filtering using advanced filter criteria, downloading link details, and resetting preferences, see "[Network Table Overview](#)" on page 515.

Tasks You Can Perform

You can perform the following tasks on the Link tab:

- From the **More** list, you can perform the following tasks:



NOTE: You can also right-click a link to view the same options.

- View details about a link—To view details such as the site name, country, site id, organization, id, and time zone:
 - Hover over a link row in the table and click the **Details** icon.
 - Select the link on the table and click **More > Show Detail**.

The **Link - Device A to Device Z** appears, which includes the Details tab.

On the Details tab, you can view the node id associated with the link, IPv4 address of the two nodes, bandwidth, live status, and the operational status of the link.

The **Link - Device A to Device Z** page is movable so that you can pin it anywhere on the screen.

- Download link-related information—Click **Download** to download detailed information about all the existing links in the topology, in CSV format.
- Add links to the network model—Click **+** to add new links. For more information see ["Create a Link \(Planner\)" on page 1352](#).

Table 332: Fields on the Link Table

Fields	Description
Name	Displays the name of the link.
Device A	Displays the name/host name of the device at which traffic enters.
Device Z	Displays the name/host name of the device at which traffic exits.

Table 332: Fields on the Link Table *(Continued)*

Fields	Description
IP A	Displays the IPv4 address of the Device A's interface, which connects to this link. Routing Director displays the IPv4 address of Device A based on Device A's active configuration. If the IPv4 addresses of Device A and Device Z are in the same subnet, Routing Director forms an interface between devices A and Z.
IP Z	Displays the IPv4 address of Device Z's interface, which connects to this link. Routing Director displays the IPv4 address of IP Z based on device Z's active configuration. If the IPv4 addresses of IP A and IP Z are in the same subnet, Routing Director forms an interface between devices A and Z.
Interface A	Displays the ingress interface name of Device A.
Interface Z	Displays the egress interface name of Device Z.
Bandwidth A	Displays the bandwidth (in bps) of the ingress link.
Bandwidth Z	Displays the bandwidth (in bps) of the egress link.
Description	Displays the description you added at the time of link creation.
TE Metric A	Displays the traffic engineering metric for Device A.
TE Metric Z	Displays the traffic engineering metric for Device B.
Admin Group A	Displays one or more admin group bits for the traffic to traverse from Device A.
Admin Group Z	Displays one or more admin group bits for the traffic to traverse from Device Z.

Table 332: Fields on the Link Table *(Continued)*

Fields	Description
RSVP BW A	Displays the RSVP bandwidth between device A and device Z.
RSVP BW Z	Displays the RSVP bandwidth between device Z and device A.
ISIS L1 A	Displays the IS-IS area level (L1) metric to reach Node Z from Node A over this link.
ISIS L1 Z	Displays the IS-IS area level (L1) metric to reach Node A from Node Z over this link.
ISIS L2 A	Displays the IS-IS area level (L2) metric to reach Node Z from Node A over this link.
ISIS L2 Z	Displays the IS-IS area level (L2) metric to reach Node A from Node Z over this link.
ISIS TE A	Displays the IS-IS TE metric to reach Node Z from Node A over this link.
ISIS TE Z	Displays the IS-IS TE metric to reach Node A from Node Z over this link.
OSPF Metric A	Displays the OSPF metric to reach Node Z from Node A over this link.
OSPF Metric Z	Displays the OSPF metric to reach Node A from Node Z over this link.
OSPF TE A	Displays the OSPF TE metric to reach Node Z from Node A over this link.
OSPF TE Z	Displays the OSPF TE to reach Node A from Node Z over this link.

RELATED DOCUMENTATION

| [Planner Workflow](#) | [1334](#)

Create a Link (Planner)

You can add a link to your offline network.

To add a link:

1. Navigate to **Planning > Networks > Offline Models > *Model-Name* > Open**.

The offline topology page is displayed with the topology map at the center and the network information table at the bottom of the page.

2. Navigate to the Link tab.
3. Click the plus icon (+).

The Create Link page appears.

4. Complete the configuration according to the guidelines in [Table 333 on page 1352](#).



NOTE: Fields marked with an asterisk (*) are mandatory.

5. Click **Submit**, to add a link.

A message indicating that the link is added appears. You are returned to the Link tab where you can view the newly added link.

Table 333: Fields on the Create Link Page

Field	Description
Name*	Specify the name of the link.
Node A*	Select the name of the ingress (source) node from the list.
Node Z*	Select the name of the egress (destination) node from the list.

Table 333: Fields on the Create Link Page *(Continued)*

Field	Description
Bandwidth A*	Enter the bandwidth (in bps) for Node A. You must enter a number immediately followed by K (Kbps), M (Mbps), or G (Gbps). For example, 10M signifying 10 Megabits per second. If you enter a value without units, bps is applied.
Bandwidth Z*	Enter the bandwidth (in bps) for Node Z. You must enter a number immediately followed by K (Kbps), M (Mbps), or G (Gbps). For example, 10M signifying 10 Megabits per second. If you enter a value without units, bps is applied.
Properties section	
Description	Enter a description for the new link that you are adding.
IPv4 A	Enter the IPv4 address of Node A.
IPv4 Z	Enter the IPv4 address of Node Z.
Interface A	Select the interface of Node A.
Interface Z	Select the interface of Node Z.
TE Metric A	Enter a metric value for Node A for use in traffic engineering (TE). Range: 1 through 16,777,215
TE Metric Z	Enter a metric value for Node Z for use in TE. Range: 1 through 16,777,215
Protocols section	

Table 333: Fields on the Create Link Page *(Continued)*

Field	Description
ISIS	Toggle the button to enable or disable IS-IS protocol. If you enable ISIS: • Enter IS-IS area level (L1, L2,) metric and TE metric for Node A • Enter IS-IS area level (L1, L2,) metric and TE metric for Node Z
OSPF	Toggle the button to enable or disable the OSPF protocol. If you enable OSPF, enter the OSPF metric and OSPF TE metric for each node.
RSVP	Toggle the button to enable or disable the RSVP protocol. If you enable RSVP, enter the RSVP bandwidth for each direction from AZ (Node A to Z) and ZA (from Z to A).
Constraints	
Admin Group A	Select the bit-level link coloring from the list for the A to Z link.
Admin Group Z	Select the bit-level link coloring from the list for the Z to A link.

RELATED DOCUMENTATION

[About the Offline Topology Page | 1343](#)

[Planner Workflow | 1334](#)

About the Tunnel Tab (Planner)

IN THIS SECTION

- [Tasks You Can Perform | 1355](#)

On the Tunnels tab of the network table (**Planning > Networks > Offline Models > *Model-Name* > Open**), view the information about label-switched paths (LSPs) in your offline network.

Tasks You Can Perform

You can perform the following tasks on the Tunnels tab:

- Download—To download tunnels-related data to your local system, click **Download**. The data is downloaded to your local system as a comma-separated values (CSV) file.
- From the More list, you can perform the following tasks:
 - View details of a tunnel—To view the details of a tunnel, you can do one of the following:
 - Select a tunnel on the table and click **More > Show Details**.
 - Hover over a tunnel row in the table and click the **Details** icon.

The **Tunnel - *Tunnel-Name*** page appears.

- Add tunnels to the network model—Click **+** to add new tunnels. For more information see ["Create a Tunnel \(Planner\)" on page 1357](#).
- Reload—To view the latest changes in the network table, click **Reload**.
- To perform common actions such as filtering using advanced filter criteria and resetting preferences, see ["Network Table Overview" on page 515](#).

Table 334: Fields on Tunnels Tab

Column	Description
Name	Displays the name of the tunnel.

Table 334: Fields on Tunnels Tab *(Continued)*

Column	Description
Device A	Displays the ingress device of the tunnel at which traffic enters.
Device Z	Displays the egress device of the tunnel at which traffic exits.
IP A	Displays the IPv4 address of Device A.
IP Z	Displays the IPv4 address of Device Z.
Bandwidth	Displays the signaled or reserved bandwidth (bps) associated with the tunnel.
Admin Group Include All	Displays one or more admin group bits for the tunnel. The tunnel traverses links that include all of the admin groups listed here. Admin groups are created to implement a variety of policy-based LSP setups.
Admin Group Include Any	Displays one or more admin group bits for the tunnel. The tunnel traverses links that include at least one of the admin groups listed here. Admin groups are created to implement a variety of policy-based LSP setups.
Admin Group Exclude	Displays one or more admin group bits for the tunnel. The tunnel traverses links that do not include any of the admin groups Admin groups are created to implement a variety of policy-based LSP setups.
Path Type	Displays the type of LSP path—primary, secondary, or standby.
Path Name	Displays the user-defined name of the LSP path.
Setup	Displays the setup priority of the LSP. The setup priority ranges from 0 (highest) through 7 (lowest).

Table 334: Fields on Tunnels Tab (*Continued*)

Column	Description
Hold	Displays the RSVP hold priority for the tunnel traffic. The hold priority ranges from 0 (highest) through 7 (lowest).
Record Route	Displays record route object (RRO), which is used to record the path that the LSP traverses. It includes the IP address and router/node ID of each RSVP node in the path.
Description	Displays the description that you added at the time of tunnel creation.

RELATED DOCUMENTATION

[Planner Workflow](#) | [1334](#)

Create a Tunnel (Planner)

Routing Director enables you to provision Label-Switched Path (LSP) in your network model. The LSPs are also known as tunnels. Create RSVP, segment-routed (SR) and SRv6 LSPs, and view the LSPs on the topology map.

To create a tunnel:

1. Navigate to **Planning > Networks > Offline Models > *Model-Name* > Open**.

The offline topology page is displayed with the topology map at the center and the network information table at the bottom of the page.

2. Navigate to the Tunnels tab.
3. Click the icon (+).

The Create Tunnel page appears.

4. Complete the configuration according to the guidelines in [Table 335 on page 1358](#).



NOTE: Fields marked with an asterisk (*) are mandatory.

5. Click **Submit**, to add a tunnel.

A message indicating that the tunnel is created appears. You are returned to the Tunnels tab where you can view the newly added tunnel.

Table 335: Fields on the Create Tunnel Page

Field	Description
Name*	Specify a unique name for the tunnel. You can use any number of alphanumeric characters, period, hyphens, colons, and underscores.
Node A*	From the list, select the ingress node of the tunnel.
Node Z*	From the list, select the egress node of the tunnel.
Provisioning Type*	From the list, select RSVP as the tunnel type.
Path Type*	From the list, select primary as the path type.
Planned Bandwidth*	Specify the planned bandwidth of the tunnel (along with valid units, with no space between the bandwidth and units) for the tunnel. If you specify a value without units, bps is automatically applied. Valid units: B or b, M or m, K or k, or G or g
Properties section	
Description	Enter a brief description about the tunnel that you are creating.

Table 335: Fields on the Create Tunnel Page *(Continued)*

Field	Description
Setup Priority	Specify the setup priority for the tunnel traffic. Priority levels range from 0 (highest priority) through 7 (lowest priority). The default is 7, which is the standard MPLS tunnel definition in Junos OS. Based on the setup priority, the PCE determines whether a new tunnel can be established, by preempting an existing tunnel. The existing tunnel can be preempted if the setup priority of the new tunnel is higher than that of the existing tunnel and the preemption releases enough bandwidth for the new tunnel.
Hold Priority	Specify the hold priority for the tunnel traffic. Priority levels range from 0 (highest priority) through 7 (lowest priority). The default is 7, which is the standard MPLS tunnel definition in Junos OS. Based on the hold priority, the PCE determines whether the tunnel can be preempted or not. If the hold priority for a tunnel is higher, it is unlikely for the tunnel to be preempted.
Routing Method	From the list, select a routing method. The routing method determines the metric that Planner uses to compute path for tunnels. Default indicates TE Metric, ISIS indicates ISIS L2 or L1 metric, and OSPF indicates OSPF metric of the links.
Constraints section	
Admin Group Exclude	Select one or more admin group bits. The tunnel traverses links that do not include any of the admin groups specified in this field. The maximum selections allowed are 32.
Admin Group Include All	Select one or more admin group bits for the tunnel to traverses links that include all of the admin groups specified in this field. The maximum selections allowed are 32.

Table 335: Fields on the Create Tunnel Page (*Continued*)

Field	Description
Admin Group Include Any	Select one or more admin group bits. The tunnel traverse links that include at least one of the admin groups specified in this field. You can assign a maximum of 32 values.

RELATED DOCUMENTATION

[About the Offline Topology Page | 1343](#)

[Planner Workflow | 1334](#)

About the Demands Tab (Planner)

IN THIS SECTION

- [Tasks You Can Perform | 1360](#)

On the Demands tab of the network table (**Planning > Networks > Offline Models > *Model-Name* > Open**), view the information about demands in your offline network.

Demands represent the traffic flow that need to be routed with specific constraints and performance goals. Demands are created when you want to analyze the traffic between the nodes.

Tasks You Can Perform

You can perform the following tasks on the Demands tab:

- **Download**—To download demands-related data to your local system, click **Download**. The data is downloaded to your local system as a comma-separated values (CSV) file.
- From the More list, you can perform the following tasks:
 - **View details of a demand**—To view the details of a demand, you can do one of the following:

- Select a demand on the table and click **More > Show Details**.
- Hover over a Demand row in the table and click the **Details** icon.

The **Demand - *Demand-Name*** page appears. You can view the To and From traffic-related information.

- Add demands to the network model—Click **+** to add new demands. For more information see ["Create a Demand \(Planner\)" on page 1362](#).
- To perform common actions such as filtering using advanced filter criteria and resetting preferences, see ["Network Table Overview" on page 515](#).

Table 336: Fields on Demands Tab

Column	Description
Name	Displays the name of the demand.
Device A	Displays the device at which demand enters.
Device Z	Displays the device at which demand exits.
IP A	Displays the IPv4 address of the interface from which device A sends traffic.
IP Z	Displays the IPv4 address of device Z interface that receives traffic.
Bandwidth	Displays the traffic rate (bps) associated with the demand.
Record Route	Displays whether Routing Director should actively record the path by sending the record route object (RRO). The RRO is used to record the path that the demand traverses. It includes the IP address and router/node ID of each RSVP node in the path.
Description	Displays the description that you added at the time of demand creation.

RELATED DOCUMENTATION

| [Planner Workflow](#) | 1334

Create a Demand (Planner)

You can create demands for your network model. Demands are created when you want to analyze the traffic between the nodes.

To create a demand:

1. Navigate to **Planning > Networks > Offline Models > *Model-Name* > Open**.

The offline topology page is displayed with the topology map at the center and the network information table at the bottom of the page.

2. Navigate to the Demands tab.

3. Click the plus icon (+).

The Create Demand page appears.

4. Complete the configuration according to the guidelines in [Table 337 on page 1362](#).



NOTE: Fields marked with an asterisk (*) are mandatory.

5. Click **Submit** to create a demand.

A message indicating that the demand is created appears. You are returned to the Demands tab where you can view the newly added demand.

Table 337: Fields on the Create Demand Page

Column	Description
Name*	Specify a unique name for the demand. You can use any number of alphanumeric characters, period, hyphens, colons, and underscores.
Node A*	From the list, select the device that you want to use as the ingress node.
Node Z*	From the list, select the device that you want to use as the egress node.

Table 337: Fields on the Create Demand Page (*Continued*)

Column	Description
Planned Bandwidth	Specify the planned bandwidth (along with valid units, with no space between the bandwidth and units) for the demand. If you specify a value without units, bps is automatically applied. Valid units: B or b, M or m, K or k, or G or g.
Properties section	
Description	Enter a brief description about the demand that you are creating.

RELATED DOCUMENTATION

[About the Offline Topology Page | 1343](#)

[Planner Workflow | 1334](#)

About the Interface Tab (Planner)

IN THIS SECTION

- [Tasks You Can Perform | 1364](#)
- [Field Descriptions | 1364](#)

On the Interfaces tab of the network table (**Planning > Networks > Offline Models > *Model-Name* > Open**), view the information about interfaces in your network model.

Use the Interfaces tab to view details of the interfaces present on the device. The details include interface name, MAC address, operational status, speed, duplex mode, and so on. You can view and modify the interface name, the IP address (IPv4 or IPv6), the administration status, and the description of the interfaces.

Tasks You Can Perform

You can perform the following tasks from this page:

- View the details of interfaces present on the chassis, line cards, FPCs, and PICs. See [Table 338 on page 1364](#).
- Add interfaces to the offline network—Click + to add new interfaces. For more information see ["Create an Interface \(Planner\)" on page 1365](#).
- To perform common actions such as filtering using advanced filter criteria and resetting preferences, see ["Network Table Overview" on page 515](#).

Field Descriptions

[Table 338 on page 1364](#) describes the fields on the Interfaces tab:

Table 338: Fields on the Interfaces Tab

Field	Description
Interface Name	Displays the name of the interface.
Device	Displays the name of the device to which the interface belongs to.
IPv4 Address/Subnet Mask	Displays the IPv4 address assigned to the interface.
Bandwidth	Displays the bandwidth of the interface.
Description	Displays the description of the interface that you have entered at the time of interface creation.
VRF	Displays the name of Virtual Routing and Forwarding (VRF) routing instance that the interface is part of.

RELATED DOCUMENTATION

Create an Interface (Planner)

On the Interface tab, add an interface for a device in your network model.

To create an interface:

1. Navigate to **Planning > Networks > Offline Models > *Model-Name* > Open**.

The offline topology page is displayed with the topology map at the center and the network information table at the bottom of the page.

2. Navigate to the Interfaces tab.

3. Click the icon (+).

The Create Interface page appears.

4. Complete the configuration according to the guidelines in [Table 339 on page 1365](#).



NOTE: Fields marked with an asterisk (*) are mandatory.

5. Click **Submit**, to create the interface.

A message indicating that the interface is created appears. You are returned to the Interface tab where you can view the newly added interface.

Table 339: Fields on the Create Interface Page

Column	Description
Name*	Specify the name of the interface. Use alphanumeric characters, underscore, hyphen, period, or colon.
Node*	Select the device for which you are configuring the interface.
Bandwidth	Specify the bandwidth allowed through the interface in bps.
Properties	
Description	Enter a description for the interface.
IPv4 Address/Subnet Mask	Specify the IPv4 address for the interface in CIDR format.

Table 339: Fields on the Create Interface Page *(Continued)*

Column	Description
VRF	Specify the name of Virtual Routing and Forwarding (VRF) routing instance that the interface belongs to.

RELATED DOCUMENTATION

About the Offline Topology Page 1343
Planner Workflow 1334

10

PART

Paragon Shell CLI Reference

- [Introduction | 1368](#)
 - [Operational Mode Commands | 1376](#)
 - [Configuration Mode Commands | 1491](#)
 - [Troubleshooting Commands | 1536](#)
 - [Service Orchestration | 1575](#)
-

Introduction

IN THIS CHAPTER

- [Paragon Shell Overview | 1368](#)
- [A Quick Tour of Paragon Shell | 1369](#)

Paragon Shell Overview

IN THIS SECTION

- [Benefits of Paragon Shell | 1369](#)

Routing Director provides a custom Containerized Management Daemon (cMGD) user shell, called Paragon Shell. The Paragon Shell CLI is text-based and similar to the Junos cMGD CLI. Paragon Shell is installed and available after the VMs are created on the VMware ESXi server using the OVA bundle. The OVA bundle is prepackaged with all the packages required to create the node VMs and deploy the Routing Director cluster. The VMs are created with a Linux base OS and Paragon Shell operates on top of the base OS. A system administrator can use Paragon Shell to deploy, configure, and monitor the Routing Director cluster.

A system administrator uses Paragon Shell to perform the following categories of tasks:

- Deploy the Routing Director cluster
- Back up and restore the cluster applications
- Upgrade the cluster
- Manage users
- Perform system reboots

- Troubleshoot the cluster and applications

In addition to Paragon Shell, you can also access the cMGD CLI for Service Orchestration tasks. To access the Service Orchestration cMGD CLI, you must exit out of Paragon Shell and explicitly deploy and log in to the Service Orchestration cMGD CLI. For more information, see ["About the Service Orchestration cMGD CLI" on page 1576](#).

Benefits of Paragon Shell

- Intuitive and easy to use text-based CLI
- Reduces the risk of unintended changes through the Linux root shell

In this chapter, you will learn about using operational commands and configuration statements. This chapter is a CLI reference of supported commands that are available in Paragon Shell and describes the syntax for their usage.

RELATED DOCUMENTATION

[A Quick Tour of Paragon Shell | 1369](#)

[About the Service Orchestration cMGD CLI | 1576](#)

A Quick Tour of Paragon Shell

IN THIS SECTION

- [CLI Modes | 1370](#)
- [Display set Commands from the Configuration | 1372](#)
- [Getting Help About Commands | 1374](#)

The following topics can help you get started with the Paragon Shell CLI to perform configuration changes, switch between operational mode and configuration mode, create a user account, and execute some of the basic commands.

CLI Modes

The CLI has two modes:

- **Operational mode**—Use this mode to display the status of and information on the Routing Director cluster including the version of the software installed. In operational mode, you enter commands to monitor and to troubleshoot the Routing Director application. When you log in to a node VM, you are placed in operational mode by default.

If you are in the Linux root shell, type `cli` to enter Paragon Shell operational mode.

```
user@node>
```

- **Configuration mode**—Use this mode to configure the cluster, SMTP, monitoring, user access, and several application properties. Type `configure` to enter configuration mode. To implement all configured changes, you must commit the changes.

```
user@node> configure
Entering configuration mode

[edit]
user@node#
```

The CLI prompt changes from `user@node>` to `user@node#`, showing that you are in configuration mode, and a banner appears to indicate the hierarchy level. You can exit configuration mode and return to operational mode in one of the following ways:

- To commit the configuration and exit:

```
[edit]
user@node# commit and-quit
commit complete
Exiting configuration mode
user@node>
```

- To exit without committing:

```
[edit]
user@node# exit
Exiting configuration mode
user@node>
```

When you exit configuration mode, the CLI prompt changes from `user@node#` to `user@node>`, and the banner no longer appears. You can enter or exit configuration mode as many times as you wish without committing your changes.

You can view and execute the commands that are available corresponding to your role and privileges. Currently, we have the `super-user` and `read-only` privileges. A superuser, or root user like the system administrator, can view and execute all the supported commands available in Paragon Shell. A user with read-only privileges can execute a limited set of commands. To view the commands available for your user role, type `?` and press enter. [Table 340 on page 1371](#) describes the commands you can use to navigate the CLI.

Table 340: Commands to Navigate Paragon Shell

Command	Description
<code>?</code>	Type <code>?</code> to view the entire list of supported commands available to you. Append <code>?</code> at the end of a command to view a list of possible options you can use to complete and execute the command. For example, request <code>paragon backup ?</code> all the options available to back up your Routing Director cluster. You can also press <code>tab</code> to auto-complete a command or show the available options for that command.
<code>cli</code>	Type this command in the Linux root shell to enter Paragon Shell.
<code>quit</code>	Exit Paragon Shell or exit from configuration mode. The <code>quit</code> and <code>exit</code> commands are equivalent.

Table 340: Commands to Navigate Paragon Shell *(Continued)*

Command	Description
exit	Type this command to exit from the Paragon Shell CLI or configuration mode. If you are in the configuration mode, type exit to exit the configuration mode to the operational mode. If you type exit in the Paragon Shell operational mode, you will exit to the Linux root shell.
configure	Type this command to enter the configuration mode in Paragon Shell.
Configuration Mode Commands	
commit	Type this command to commit any configuration updates you make in the configuration mode.
commit and quit	Type this command to commit any configuration updates you make and exit the configuration mode.
set	Add or edit a configuration.
show show <i>statement-path</i>	Display the current configuration.
delete	Delete a configuration.

Display set Commands from the Configuration

In configuration mode, you can display the configuration as a series of configuration mode commands required to re-create the configuration. This is useful if you are not familiar with how to use configuration mode commands or if you want to cut, paste, and edit the displayed configuration.

To display the configuration as a series of configuration mode commands, which are required to re-create the configuration from the top level of the hierarchy as set commands, issue the `show configuration` mode command with the `display set` option:

```
user@node# show | display set
```

For example:

```
user@node# show | display set
set version "20240408.195519__cd-builder.r1414545 [_cd-builder]"
set system root-authentication encrypted-password "$6$lwye5$RyPtoSkdX9.X10sxPFt0lJ2g/
A8EYU4vj26Lq6TslH6yH0uytvVg93bJK05KC8GQI88Q..mCug7zG0mcaxIBd/"
set system scripts commit file sync_on_commit.py
set system scripts action max-datasize 2g
set system login user name uid 2001
set system login user name class super-user
set system login user name authentication encrypted-password "\$6\$LitsWvUBaJSowtHo\$BK5ze3gCvMS/
ZhIGHFpwHv1QuNpCUy0aJxjrk5Zsm8xGzs6jfeiBKs7rX0cYG.KMCmHqpBIw.rLunX6xahH8I."
set system login user paragonadmin uid 2000
set system login user paragonadmin class super-user
set system login user paragonadmin authentication encrypted-password "\$6\$oMaQQTDLvno8D9.Q\
$dErpR2riP9yu3Ry1H.MKRnSlMsc93fR04S4QRx.C9NhnZci8sKi1AyU0nYpkfW0V2LIJJYHAAawLwKENusUw.."
set system services ssh port 2222
set system services extension-service request-response grpc clear-text port 50051

<output snipped>
```

To display the configuration as set commands and search for text matching a regular expression by filtering output, specify the `match` option after the pipe (`|`):

```
user@node# show | display set | match regular-expression
```

For example:

```
uaer@node# show | display set | match user
set system login user name uid 2001
set system login user name class super-user
set system login user name authentication encrypted-password "\$6\$LitsWvUBaJSowtHo\$BK5ze3gCvMS/
ZhIGHFpwHv1QuNpCUy0aJxjrk5Zsm8xGzs6jfeiBKs7rX0cYG.KMCmHqpBIw.rLunX6xahH8I."
set system login user paragonadmin uid 2000
```

```
set system login user paragonadmin class super-user
set system login user paragonadmin authentication encrypted-password "$6$0MaQQTDLvno8D9.Q\
$dErpR2riP9yu3Ry1H.MKRnSlMsc93fR04S4QRx.C9NhnZci8sKi1AyU0nYpkfW0V2L1JJYHAAawLwKENusUw.."
set paragon cluster applications web-ui web-admin-user "name@juniper.net"
```

```
[edit]
```

Getting Help About Commands

CLI commands and options can vary by software release, your role, and privileges. Each level of the CLI command hierarchy provides information about available commands. You can type a question mark (?) to get context-relevant help about commands.

If you type the question mark at the command-line prompt, the CLI lists the available commands and options. For example, to view a list of top-level operational mode commands, this is the result:

```
user@node> ?
Possible completions:
  clear          Clear information in the system
  configure      Manipulate software configuration information
  file           Perform file operations
  monitor        Show real-time debugging information
  quit           Exit the management session
  request        Make system-level requests
  set            Set CLI properties, date/time, craft interface message
  show           Show system information
  ssh            Start secure shell on another host
user@node>
```

If you type the question mark after entering the complete name of a command or command option, the CLI lists the available commands and options and then re-displays the command names and options you typed.

```
user@node> request ?
Possible completions:
  paragon
  system          Perform system-level operations
```

If you type the question mark in the middle of a command name, the CLI lists possible command completions that match the letters you have entered so far. It then re-displays the letters that you typed. For example, to list all operational mode commands that start with the letter **s**, type the following:

```
user@node> s?  
Possible completions:  
  set          Set CLI properties, date/time, craft interface message  
  show         Show system information  
  ssh          Start secure shell on another host
```


Operational Mode Commands

IN THIS CHAPTER

- file copy | 1377
- monitor | 1379
- request paragon backup | 1381
- request paragon cluster pods reset | 1389
- request paragon cluster upgrade | 1391
- request paragon config | 1398
- request paragon deploy | 1400
- request paragon deploy cluster | 1403
- request paragon destroy cluster input | 1406
- request paragon fix-permission | 1408
- request paragon health-check | 1409
- request paragon load | 1412
- request paragon repair-node | 1413
- request paragon replace-node | 1414
- request paragon restore | 1416
- request paragon reboot | 1418
- request paragon running-config | 1420
- request paragon shutdown | 1422
- request paragon service | 1424
- request paragon ssh | 1426
- request paragon ssh-key | 1427
- request paragon storage cleanup | 1431
- request paragon super-user password reset | 1433
- request system decrypt password | 1435
- request system reboot | 1436
- request paragon trust data update hostname | 1438

- [show configuration paragon cluster | 1439](#)
- [show configuration paragon cluster applications \(o\) | 1448](#)
- [show configuration paragon cluster common-services \(o\) | 1451](#)
- [show host disk usage | 1454](#)
- [show paragon backup | 1457](#)
- [show paragon certificate expiry-date certificate-type | 1462](#)
- [show paragon cluster \(o\) | 1465](#)
- [show paragon cluster details | 1467](#)
- [show paragon cluster namespaces | 1468](#)
- [show paragon cluster nodes | 1470](#)
- [show paragon cluster pods | 1473](#)
- [show paragon cluster pods namespace healthbot sort | 1476](#)
- [show paragon images version | 1479](#)
- [show paragon images version namespace | 1481](#)
- [show paragon pvc details | 1484](#)
- [show paragon version | 1489](#)

file copy

IN THIS SECTION

- [Syntax | 1378](#)
- [Description | 1378](#)
- [Options | 1378](#)
- [Required Privilege Level | 1378](#)
- [Sample Output | 1378](#)
- [Release Information | 1379](#)

Syntax

```
file copy source-path destination-path
```

Description

Copy a file from one location to another.

Options

<i>source-path</i>	Source file and path.
<i>destination-path</i>	Destination path where the file must be copied.

Required Privilege Level

configure

Sample Output

```
file copy /var/log/action.log root@10.1.2.3:/tmp
```

```
user1@node> file copy /var/log/action.log root@10.1.2.3:/tmp
```

Where:

- *source-path:/var/log/action.log* is the file path inside the Docker container of node, in which user1 is logged in.
- *destination-path:/tmp* is the file path in node 10.1.2.3 but not inside any Docker container as no user is logged into it.

```
file copy root@10.1.2.3:/tmp/action1.log /var/log/
```

```
user1@node> file copy root@10.1.2.3:/tmp/action1.log /var/log/
```

Where:

- *source-path:/tmp/action1.log* in node 10.1.2.3 but not inside any Docker container as no user is logged into it.
- *destination-path:/var/log/* is the file path inside the Docker container of node, in which user1 is logged in.

Release Information

Command introduced in Juniper Paragon Automation Release 2.0.0.

monitor

IN THIS SECTION

- [Syntax | 1379](#)
- [Description | 1379](#)
- [Options | 1380](#)
- [Required Privilege Level | 1380](#)
- [Output Fields | 1380](#)
- [Sample Output | 1380](#)
- [Release Information | 1381](#)

Syntax

```
monitor (start|stop) log_filename
```

```
monitor list
```

Description

Show real-time debugging information.

Options

<code>start <i>log_filename</i></code>	Start displaying a log file and additional entries being added to the log file in real time.
<code>stop <i>log_filename</i></code>	Stop displaying a log file in real time.
<code>list</code>	Display a list of monitored files.

Required Privilege Level

configure

Output Fields

When you enter this command, you are provided feedback on the status of your request or you are returned to the command prompt.

Sample Output

`monitor start filename`

```
user@node> monitor start /epic/config/log  
  
user@node>
```

`monitor stop filename`

```
user@node> monitor stop /epic/config/log  
  
user@node>
```

monitor list

```
user@node> monitor list  
monitor start "/epic/config/log" (Last changed Apr 16 18:02:13)
```

Release Information

Command introduced in Juniper Paragon Automation Release 2.0.0.

request paragon backup

IN THIS SECTION

- [Syntax | 1382](#)
- [Description | 1382](#)
- [Options | 1383](#)
- [Required Privilege Level | 1385](#)
- [Output Fields | 1385](#)
- [Sample Output | 1386](#)
- [Release Information | 1388](#)

Syntax

```
request paragon backup
```

```
request paragon backup config
```

```
request paragon backup delete backup-id backupID
```

```
request paragon backup download backup-id backupID storage-location remote-path user username  
password password
```

```
request paragon backup telemetry
```

```
request paragon backup upload backup-id backupID storage-location remote-path user username  
password password
```

```
request paragon backup schedule trigger-time schedule storage-location remote-path user username  
password password node nodeIP type (full | config | telemetry)
```

Description

Back up your Routing Director network configuration and telemetry data. Configuration data includes information stored in PostgreSQL, ArangoDB, and Airflow configuration database systems, along with all available software images, and AskParagon secrets keys. Telemetry data includes information stored in database systems such as OpenSearch, TimescaleDB including Active Assurance data, and VictoriaMetrics including Insights and Pathfinder data.

Additionally, you can schedule backups, upload your backup file to a remote location outside the Routing Director cluster and download a backup file from a remote location. You can also delete the backed-up data.

Options

none	Back up your current Routing Director network configuration and telemetry data. Backup files are available at export/paragon-shell/backup .
config	Back up only your current Routing Director network configuration information.
delete backup-id <i>backupID</i>	Delete the backed-up file with the specified <i>backupID</i> . You can delete only those files present on the node that you are logged on to. To view a list of available backup files, see " show paragon backup " on page 1457.
download backup-id <i>backupID</i> storage-location <i>path</i> user <i>username</i> password <i>password</i>	Download a backed-up file from a remote location. Download options include: • backup-id <i>backupID</i>—ID of the backup file. • storage-location <i>path</i>—Path where the backup file is stored remotely. Enter the location in the following format: scp://host:port/path or sftp://host:port/path • user <i>username</i>—Username to log in to the remote location. • (Optional) password <i>password</i>—Password to log in to the remote location. If you do not enter the password, you will be prompted to enter it.
telemetry	Back up only your current Routing Director telemetry information.

<code>upload backup-id <i>backupID</i> storage-location <i>path</i> user <i>username</i> password <i>password</i></code>	Upload the backup files to a remote location. Upload options include: • backup-id <i>backupID</i>—ID of the backup file. • storage-location <i>path</i>—Path where the backup files must be stored remotely. Enter the location in the following format: scp://host:port/path or sftp://host:port/path • user <i>username</i>—Username to log in to the remote location. • password <i>password</i>—Password to log in to the remote location. If you do not enter the password, you will be prompted to enter it.

schedule trigger-time <i>schedule</i> node <i>nodeIP</i> storage-location <i>remote-path</i> user <i>username</i> password <i>password</i> type (<i>full</i> / <i>config</i> / <i>telemetry</i>)	Schedule periodic backups. Scheduling options include: • trigger-time <i>schedule</i>—The cron expression of the backup schedule. • node <i>hostname</i>—Hostname of the node on which you want to execute the backup scheduling command and store the backup job logs. Use the <code>kubect1 get nodes</code> command on the Linux root shell to determine the correct hostname. • storage-location <i>remote-path</i>—Path where the backup files must be stored remotely. Enter the location in the following format: scp://host:port/path or sftp://host:port/path • user <i>username</i>—Username to log in to the remote location. • password <i>password</i>—Password to log in to the remote location. If you do not enter the password, you will be prompted to enter it. • type (<i>full</i> / <i>config</i> / <i>telemetry</i>)—Back up your current Routing Director network configuration information, telemetry information, or both. If not specified, both configuration and telemetry information are backed up. Rerun this command to edit an existing backup schedule.

Required Privilege Level

configure

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request paragon backup

```

user@node> request paragon backup
Postgres backup started
Postgres backup completed
Using arango CRDN pod foghorn-crdn-drjaslb3-413778
ArangoDB backup completed
Using workflow manager pod workflow-manager-64f666454c-fmc62
Airflow backup completed
Checking available software images
Downloading software images
Image download progress :  0%|          | 0/14 [00:00]
Image download progress :  7%|7        | 1/14 [00:17]
Image download progress : 14%|#4       | 2/14 [00:35]
Image download progress : 29%|##8      | 4/14 [00:35]
Image download progress : 50%|#####   | 7/14 [00:35]
Image download progress : 71%|#####1  | 10/14 [00:35]
Image download progress : 93%|#####2| 13/14 [00:36]
Image download progress : 100%|#####| 14/14 [00:36]
Software image backup completed
AskParagon backup completed
Timescaledb backup job started
Opensearch backup job started
insights Victorimetrics backup job started
pathfinder Victorimetrics backup job started
Backup job in progress for 20250401-024304

```

request paragon backup config

```

user@node> request paragon backup config
Postgres backup started
Postgres backup completed
Using arango CRDN pod foghorn-crdn-o0b3riuv-413778
ArangoDB backup completed
Airflow backup completed
Checking available software images
No software images available for backup

```

```
AskParagon backup completed
Backup completed for 20250305-175217
```

request paragon backup delete backup-id *backup-id*

```
user@node> request paragon backup delete backup-id 20240409-102055
Successfully deleted backup 20241202-192055
```

request paragon backup download

```
user@node> request paragon backup download backup-id 20241111-050227 storage-location scp://
10.1.2.3/root/storage user root password LoginPassword
Downloading backup for 20241111-050227 is in-progress
```

request paragon backup telemetry

```
user@node> request paragon backup telemetry
Timescaledb backup job started
Opensearch backup job started
insights Victoriameetrics backup job started
pathfinder Victoriameetrics backup job started
Backup job in progress for 20250305-175259
```

request paragon backup upload

```
user@node> request paragon backup upload backup-id 20241111-050227 storage-location scp://
10.1.2.3/root/storage user root password LoginPassword
Uploading backup for 20241111-050227 is in-progress
```

request paragon backup schedule

```
root@node> request paragon backup schedule storage-location scp://10.1.2.3/tmp/backup
Provide backup scheduling time
How often should the backup run?
1) Hourly
```

```

2) Daily
3) Weekly
4) Monthly
5) Custom
Select an option (1-5): 5
Minute (0-59) [0]:
Hour (0-23) [0]:
Day of month (1-31) [*]: 1
Month (1-12) [*]: */3
Day of week (0-6) [*]:
Trig Time 0 0 1 */3 *

Storage User: root
Storage Password:
Backup scheduled : At 12:00 AM, on day 1 of the month, every 3 months

```

request paragon backup delete schedule

```

root@node> request paragon backup delete schedule
Removed scheduled backup

```

Release Information

Command introduced in Juniper Paragon Automation Release 2.0.0.

config, telemetry, upload, and download options introduced in Juniper Paragon Automation Release 2.3.0.

trigger-time and delete schedule options introduced in Juniper Routing Director Release 2.5.0.

RELATED DOCUMENTATION

[show paragon backup](#) | 1457

[request paragon restore](#) | 1416

request paragon cluster pods reset

IN THIS SECTION

- [Syntax | 1389](#)
- [Description | 1389](#)
- [Options | 1389](#)
- [Required Privilege Level | 1390](#)
- [Output Fields | 1390](#)
- [Sample Output | 1390](#)
- [Release Information | 1390](#)

Syntax

```
request paragon cluster pods reset namespace operation (restart|start|stop) namespace-name
service service-name
```

Description

Start, stop, or restart a service within a cluster namespace. You can also use this command to temporarily stop a service to apply new configuration changes or to fix memory issues, restart a service to apply new deployment changes, restart a stopped service, and so on.

Options

operation (restart start stop)	(Mandatory) Restart, start, or stop the service.
namespace namespace-name	(Mandatory) Specify the cluster namespace with the required service.
service service-name	(Mandatory) Perform the required operation on the service.

Required Privilege Level

configure

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request paragon cluster pods reset operation restart

```
user@node> request paragon cluster pods reset operation restart namespace healthbot service tand
Replica count for healthbot, tand in deployment is: 3
Replica count for healthbot, tand is set to: 0 in deployment, deployment.apps/tand scaled
Replica count for healthbot, tand is set to: 3 in deployment, deployment.apps/tand scaled
The namespace and service healthbot, tand has been restarted!
```

request paragon cluster pods reset operation start

```
user@node> request paragon cluster pods reset operation start namespace healthbot service tand
Replica count for healthbot, tand in deployment is: 3
healthbot, tand is already running and set to 3. To stop it, please use operation: stop.
```

request paragon cluster pods reset operation stop

```
user@node> request paragon cluster pods reset operation stop namespace healthbot service tand
Replica count for healthbot, tand in deployment is: 3
Replica count for healthbot, tand is set to 3 in configmap, configmap/replicas-configmap patched
Replica count for healthbot, tand is set to: 0 in deployment, deployment.apps/tand scaled
The namespace and service healthbot, tand has been stopped!
```

Release Information

Command introduced in Juniper Paragon Automation Release 2.0.0.

request paragon cluster upgrade

IN THIS SECTION

- [Syntax | 1391](#)
- [Description | 1391](#)
- [Options | 1391](#)
- [Required Privilege Level | 1394](#)
- [Output Fields | 1394](#)
- [Sample Output | 1395](#)
- [Release Information | 1398](#)

Syntax

```
request paragon cluster upgrade local filename upgrade-filename (none|no-confirm|detach-process|
input input-string|patch)
```

```
request paragon cluster upgrade remote url web-url (none|no-confirm|detach-process|disk-saving|
input input-string|patch)
```

Description

Upgrade your Routing Director cluster. You can upgrade the cluster by downloading the upgrade file locally to your cluster node or by using the upgrade file directly from a remote server.

Options

local filename <i>upgrade-filename</i>	Upgrade the Routing Director cluster using a local upgrade file downloaded to the /root/epic/temp/ directory.

local filename <i>upgrade-filename</i> no-confirm	Upgrade the Routing Director cluster using a local file and ignore the cluster health-check AMBER status and continue with upgrade. You are not prompted to confirm or stop or continue with the upgrade.
local filename <i>upgrade-filename</i> detach-process	The upgrade process runs in the background and frees up the CLI screen for any other tasks. The command runs the initial health checks and then proceeds with the upgrade. Once the upgrade process starts, the process is detached and moved into the background and you are returned to the command prompt. The upgrade output is logged in the /epic/temp/upgrade.log file. To monitor the status of the upgrade process and print the output onscreen, use the monitor start <code>/epic/temp/upgrade.log</code> command. When the upgrade process completes, a success message similar to the following is displayed on all the cluster nodes: Paragon Cluster upgrade is successful! - Run 'request paragon health-check' command to check current system health with upgraded Paragon cluster. - Please continue to primary host node to upgrade Paragon-shell and update OVA system files by: <code>/root/epic/upgrade_paragon-shell_ova-system.sh</code> If you get disconnected from the VM during the upgrade process, you can periodically check the upgrade log file for status on the upgrade.
local filename <i>upgrade-filename</i> input <i>input-string</i>	Use the input <i>input-string</i> options to pass additional Ansible input parameters to the upgrade command. For example, to enable verbose logging: local filename <i>upgrade-filename</i> input "-v"

local filename <i>upgrade-filename</i> patch	Apply a patch file to upgrade the Routing Director cluster to the patch version. The patch process is shorter than the upgrade process and patch files are smaller in size. NOTE: Do not use the patch option to upgrade the cluster using a regular upgrade.tgz file.
remote url <i>web-url</i>	Upgrade the cluster using an upgrade file located on a remote site.
remote url <i>web-url</i> /no-confirm	Upgrade the Routing Director cluster using a remote file and ignore the cluster health-check AMBER status and continue with upgrade. You are not prompted to confirm or stop or continue with the upgrade.
remote url <i>web-url</i> /detach-process	The upgrade process runs in the background and frees up the CLI screen for any other tasks. The command runs the initial health checks and then proceeds with the upgrade. Once the upgrade process starts, the process is detached and moved into the background and you are returned to the command prompt. The upgrade output is logged in the /epic/temp/upgrade.log file. To monitor the status of the upgrade process and print the output onscreen, use the monitor start /epic/temp/upgrade.log command. When the upgrade process completes, a success message similar to the following is displayed on all the cluster nodes: Paragon Cluster upgrade is successful! - Run 'request paragon health-check' command to check current system health with upgraded Paragon cluster. - Please continue to primary host node to upgrade Paragon-shell and update OVA system files by: /root/epic/upgrade_paragon-shell_ova-system.sh If you get disconnected from the VM during the upgrade process, you can periodically check the upgrade log file for status on the upgrade.

remote url <i>web-url</i> disk-saving	Delete the <i>upgrade.tgz</i> file as soon as it is unzipped. The upgrade command extracts the contents of the <i>upgrade.tgz</i> file at the beginning of the upgrade process. This option deletes the zipped file as soon as it is unzipped to free up space on the primary node. The advantage of using this option is that you need lesser free space for the upgrade process. The default minimum free space required is 15% of the total disk space + three times the upgrade file size. With this option you need a minimum free space of 15% of the total disk space + two times the upgrade file size. The disadvantage of using this command is that if upgrade fails or is interrupted, you need to re-download the <i>upgrade.tgz</i> file again.
remote url <i>web-url</i> input <i>input-string</i>	Use the input <i>input-string</i> options to pass additional Ansible input parameters to the upgrade command. For example, to enable verbose logging: <pre>remote url <i>web-url</i> input "-v"</pre>
remote url <i>web-url</i> patch	Apply a patch file to upgrade the Routing Director cluster to the patch version. The patch process is shorter than the upgrade process and patch files are smaller in size. NOTE: Do not use the patch option to upgrade the cluster using a regular upgrade.tgz file.

Required Privilege Level

configure

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request paragon cluster upgrade local filename *upgrade-filename*

```
root@node> request paragon cluster upgrade local filename upgrade_paragon-
release-2.3.0.6688.g63026f4af7.tgz
Checking paragon cluster system health before proceeding with cluster upgrade. This will take a
minute...
...
<output snipped>
...
=====
Overall cluster status
=====

GREEN

=====
Paragon cluster is healthy.
Proceed with Paragon cluster upgrade.
Upgrade is in progress ...
Updated to build: paragon-release-2.3.0.6688.g63026f4af7
Paragon Cluster upgrade is successful!
Run 'request paragon health-check' command to check current system health with upgraded Paragon
cluster.
Please continue to primary host node to upgrade Paragon-shell and update OVA system files by:
/root/epic/upgrade_paragon-shell_ova-system.sh
```

request paragon cluster upgrade local filename *upgrade-filename*

```
root@node> request paragon cluster upgrade local filename upgrade_paragon-
release-2.3.0.6688.g63026f4af7.tgz
Checking paragon cluster system health before proceeding with cluster upgrade. This will take a
minute...
...
<output snipped>
...
=====
Overall cluster status
=====
```

RED

=====

[Error] Paragon cluster is in bad state. Please fix it before trying to upgrade.

request paragon cluster upgrade local filename *upgrade-filename* detach-process

```
root@node> request paragon cluster upgrade local filename release-2.4.0.6234.g63026f4af7.tgz
detach-process
Checking paragon cluster system health before proceeding with cluster upgrade. This will take
a minute...
...
<output snipped>
...
=====
Overall cluster status
=====
GREEN
=====
Paragon cluster is healthy.
Proceed with Paragon cluster upgrade.
Upgrade is in progress ...
Paragon Cluster upgrade process is detached and running in the background...
Run 'monitor start /epic/temp/upgrade.log' command to watch upgrade progress!
```

request paragon cluster upgrade remote url *web-url*

```
root@node> request paragon cluster upgrade remote url https://10.2.3.4/paragon-images/
upgrade_paragon-release-2.3.0.6688.g63026f4af7.tgz
Checking paragon cluster system health before proceeding with cluster upgrade. This will take a
minute...
...
<output snipped>
...
=====
Overall cluster status
=====
```

AMBER

=====

[Warning] Paragon cluster is unhealthy.

Do you want to proceed? (Y/N) y

Proceed with cluster upgrade.

Upgrading paragon cluster from https://10.2.3.4/paragon-images/

Downloading tarball file upgrade_paragon-release-2.3.0.6688.g63026f4af7

Download file size: 28,831,677,064 bytes

Current disk Usage:

Total: 263,622,004,736 bytes

Used: 106,109,399,040 bytes

Available: 145,685,159,936 bytes

Please wait for current download to finish... (File is large. It may take a while.)

Upgrade tarball file is downloaded.

Upgrade is in progress ...

Updated to build: paragon-release-2.3.0.6688.g63026f4af7

Paragon Cluster upgrade is successful!

Run 'request paragon health-check' command to check current system health with upgraded

Paragoncluster.

Please continue to primary host node to upgrade Paragon-shell and update OVA system files by:

/root/epic/upgrade_paragon-shell_ova-system.sh

request paragon cluster upgrade remote url *web-url*/no-confirm

```
root@node> request paragon cluster upgrade remote url https://10.2.3.4/paragon-images/
```

```
upgrade_paragon-release-2.3.0.6688.g63026f4af7.tgz no-confirm
```

Checking paragon cluster system health before proceeding with cluster upgrade. This will take a minute...

...

<output snipped>

...

=====

Overall cluster status

=====

AMBER

=====

```
[Warning] Paragon cluster is unhealthy.
Proceed with cluster upgrade with no-confirm flag being set.
Upgrading paragon cluster from https://10.2.3.4/paragon-images/
Downloading tarball file      upgrade_paragon-release-2.3.0.6688.g63026f4af7
Download file size: 28,831,677,064 bytes
Current disk Usage:
    Total: 263,622,004,736 bytes
    Used: 106,109,399,040 bytes
    Available: 145,685,159,936 bytes
Please wait for current download to finish... (File is large. It may take a while.)
Upgrade tarball file is downloaded.
Upgrade is in progress ...
Updated to build: paragon-release-2.3.0.6688.g63026f4af7
Paragon Cluster upgrade is successful!
Run 'request paragon health-check' command to check current system health with upgraded Paragon
cluster.
Please continue to primary host node to upgrade Paragon-shell and update OVA system files by:
/root/epic/upgrade_paragon-shell_ova-system.sh
```

Release Information

Command introduced in Juniper Paragon Automation Release 2.0.0.

remote url keyword introduced in Juniper Paragon Automation Release 2.3.0.

no-confirm option introduced in Juniper Paragon Automation Release 2.3.0.

detach-process, disk-saving, input, and patch options introduced in Juniper Paragon Automation Release 2.4.0 .

request paragon config

IN THIS SECTION

- [Syntax | 1399](#)
- [Description | 1399](#)
- [Options | 1399](#)
- [Required Privilege Level | 1399](#)

- [Output Fields | 1399](#)
- [Sample Output | 1399](#)
- [Release Information | 1399](#)

Syntax

```
request paragon config
```

Description

Generate the Routing Director cluster inventory and configuration files. The **inventory** and **config.yml** files are generated and saved in the **/epic/config/** folder.

Options

This command has no options.

Required Privilege Level

configure

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

```
user@node> request paragon config
Paragon inventory file saved at /epic/config/inventory
Paragon config file saved at /epic/config/config.yml
```

Release Information

Command introduced in Juniper Paragon Automation Release 2.0.0.

request paragon deploy

IN THIS SECTION

- [Syntax | 1400](#)
- [Description | 1400](#)
- [Options | 1400](#)
- [Required Privilege Level | 1401](#)
- [Output Fields | 1401](#)
- [Sample Output | 1401](#)
- [Release Information | 1402](#)

Syntax

```
request paragon deploy cluster input input
```

```
request paragon deploy monitoring
```

```
request paragon deploy user
```

```
request paragon deploy user recover (false | true)
```

Description

Deploy the cluster and Paragon Shell user changes, and monitor the cluster.

Options

cluster	Deploy your Routing Director cluster. See "request paragon deploy cluster" on page 1403.
monitoring	Deploy the cluster monitoring changes.
user	Deploy the last committed Paragon Shell user changes.
user recover (false true)	When true, recover and add all current configured Paragon Shell users to Paragon Shell only on this node. Use when a Paragon Shell user is not configured on this node.

Required Privilege Level

configure

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request paragon deploy user

```

user@node> request paragon deploy user
Getting user-config configmap...
Added user: ['username ']
Modified user: []
Deleted user: []
warning: *** operating on 10.1.2.6 ***
warning: *** operation on 10.1.2.6 succeeds ***
warning: *** operating on 10.1.2.4 ***
warning: *** operation on 10.1.2.4 succeeds ***
warning: *** operating on 10.1.2.3 ***
warning: *** operation on 10.1.2.3 succeeds ***
warning: *** operating on 10.1.2.5 ***
warning: *** operation on 10.1.2.5 succeeds ***

```

```
Deleting user-config configmap...
Creating new user-config configmap...
```

request paragon deploy monitoring

```
user@node> request paragon deploy monitoring
Getting vector daemonset metadata...
Loading vector sources and sinks...
Validating config...
Deleting existing vector configmap...
Creating new vector configmap...
configmap/vector-config created
Vector source or sinks missing... Suppressing vector pods
```

request paragon deploy user recover true

```
root@node> request paragon deploy user
Getting user-config configmap...
Added user: ['pounds']
Modified user: []
Deleted user: []
operation succeeds on dv-primary2
operation succeeds on dv-primary1
operation succeeds on dv-primary3
operation succeeds on dv-worker1
Deleting user-config configmap...
Creating new user-config configmap...
```

Release Information

Command introduced in Juniper Paragon Automation Release 2.0.0.

RELATED DOCUMENTATION

request paragon deploy cluster 1403
request paragon destroy cluster input 1406

| [set paragon monitoring](#) | 1519

request paragon deploy cluster

IN THIS SECTION

- [Syntax](#) | 1404
- [Description](#) | 1404
- [Options](#) | 1404
- [Required Privilege Level](#) | 1405
- [Output Fields](#) | 1405
- [Sample Output](#) | 1405
- [Release Information](#) | 1406

Syntax

```
request paragon deploy cluster
```

```
request paragon deploy cluster input "-v -t apps"
```

```
request paragon deploy cluster input "-t ingress-controller"
```

```
request paragon deploy cluster input "-t addon-apps,papi -e target_components=papi-kafka-emitter,mailservice"
```

```
request paragon deploy cluster input "-t metallb,addon-apps -e target_components=pceserver"
```

```
request paragon deploy cluster input "-v -t victoriametrics-cluster"
```

Description

Deploy the Routing Director cluster as a whole or deploy individual components and applications.

Before deploying the cluster or redeploying specific applications, you must run the `request paragon config` command to update the **inventory** and **config.yml** files.

Options

none	Deploy the Routing Director cluster.
input "-v -t apps"	Redeploy all the Routing Director components and applications after a restore operation, in verbose mode.
input "-t apps"	Redeploy all the Routing Director components and applications after a restore operation.

input "-t ingress-controller"	Deploy custom user certificates.
input "-t addon-apps,papi -e target_components=papi-kafka-emitter,mailservice"	Deploy the SMTP configuration.
input "-t metallb,addon-apps -e target_components=pceserver"	Deploy the PCE Server application.
input "-t ems"	Deploy the EMS application.
input "-t healthbot"	Deploy the Paragon Insights application.
input "-v -t victoriametrics-cluster"	Deploy the Victoria Metrics cluster.

Required Privilege Level

configure

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request paragon deploy cluster

```
user@node> request paragon deploy cluster
Process running with PID: 2314003
To track progress, run 'monitor start /epic/config/log'
```

request paragon deploy cluster input *input*

```
user@node> request paragon deploy cluster input "-v -t ems -t healthbot"
Process running with PID: 2314012
To track progress, run 'monitor start /epic/config/log'
```

Release Information

Command introduced in Juniper Paragon Automation Release 2.0.0.

RELATED DOCUMENTATION

[request paragon restore | 1416](#)

[request paragon service destroy | 1425](#)

request paragon destroy cluster input

IN THIS SECTION

- [Syntax | 1406](#)
- [Description | 1406](#)
- [Options | 1406](#)
- [Required Privilege Level | 1407](#)
- [Output Fields | 1407](#)
- [Sample Output | 1407](#)
- [Release Information | 1407](#)

Syntax

```
request paragon destroy cluster input "-v -t apps"
```

Description

Destroy the Routing Director cluster applications and services before a restore operation.

Options

input "-v -t apps"	Stop all components in verbose mode.
input "-t apps"	Stop all components.

Required Privilege Level

configure

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request paragon destroy cluster -v -t apps

```
user@node> request paragon destroy cluster input "-v -t ems -t healthbot"
Process running with PID: 2314054
To track progress, run 'monitor start /epic/config/log'
```

Release Information

Command introduced in Juniper Paragon Automation Release 2.0.0.

RELATED DOCUMENTATION

request paragon restore 1416
request paragon deploy cluster 1403
request paragon service 1424

request paragon fix-permission

IN THIS SECTION

- [Syntax | 1408](#)
- [Description | 1408](#)
- [Options | 1408](#)
- [Required Privilege Level | 1409](#)
- [Sample Output | 1409](#)
- [Release Information | 1409](#)

Syntax

```
request paragon fix-permission
```

```
request paragon fix-permission password password
```

```
request paragon fix-permission username username
```

Description

Fix login permissions for the **config** directory. In certain scenarios, if the permissions to the **config** directory has changed, you must fix permissions before being able to rerun deployment.

Options

none	Fix login permissions for the config directory.
password <i>password</i>	Login password for the cluster node.

username <i>username</i>	Login username for the cluster node.

Required Privilege Level

configure

Sample Output

request paragon fix-permission

```
user@node> request paragon fix-permission
Please enter SSH username for the localhost: root
Please enter SSH password for the localhost: password
Permissions fixed successfully
```

Release Information

Command introduced in Juniper Paragon Automation Release 2.0.0.

request paragon health-check

IN THIS SECTION

- [Syntax | 1410](#)
- [Description | 1410](#)
- [Options | 1410](#)
- [Required Privilege Level | 1410](#)
- [Output Fields | 1410](#)
- [Sample Output | 1410](#)
- [Release Information | 1411](#)

Syntax

```
request paragon health-check
```

Description

Retrieve and display the status of the Routing Director cluster health. The health-check command checks for multiple parameters such as the status of the node (readiness, diskpressure, pod-status, CPU and memory usage, taints-status, and I/O latency), status of databases and services, and so on.

The cluster status health includes green, amber, and red. A green status indicates a healthy cluster, and a red status indicates serious issues in the cluster. An amber status indicates that there may be certain noncritical issues in the cluster. The status is returned amber in the following instances:

- Nodes have taints
- Disk usage or memory usage on any node exceeds 80% of available space.
- Disk I/O latency on any node exceeds 100000 ms
- Rook ceph status shows HEALTH_WARN

Options

This command has no options.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request paragon health-check

```
user@node1> request paragon health-check
Checking health on node: node1
Please wait as execution takes sometime. Health check in process.....
```

```

Health status checking...
=====
Get node count of Kubernetes cluster.
=====
OK
There are 4 nodes in the cluster.
=====
Get node status of Kubernetes cluster.
=====
OK
nodes are in the Ready state.
NAME      STATUS    ROLES                  AGE      VERSION
node1     Ready     control-plane,etcd,master 5d20h    v1.31.1+rke2r1
node2     Ready     control-plane,etcd,master 5d20h    v1.31.1+rke2r1
node3     Ready     control-plane,etcd,master 5d20h    v1.31.1+rke2r1
node4     Ready     worker                    5d20h    v1.31.1+rke2r1
...
<output snipped>
...
=====
Verifying Healthbot Victoria Metrics
=====
OK
=====
Verifying Northstar Victoria Metrics
=====
OK
=====
Overall cluster status
=====
GREEN

```

Release Information

Command introduced in Juniper Paragon Automation Release 2.3.0.

request paragon load

IN THIS SECTION

- [Syntax | 1412](#)
- [Description | 1412](#)
- [Options | 1412](#)
- [Required Privilege Level | 1412](#)
- [Output Fields | 1413](#)
- [Sample Output | 1413](#)
- [Release Information | 1413](#)

Syntax

```
request paragon load
```

Description

Load user configuration to the CLI.

Options

monitoring	For internal use only.
user	Load CLI User configuration into CLI.

Required Privilege Level

configure

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

```
user@node> request paragon load monitoring
Getting vector daemonset namespace...
Getting vector configmap...
Translating vector config...
Loading paragon monitoring config to mgd...
```

Release Information

Command introduced in Juniper Paragon Automation Release 2.0.0.

request paragon repair-node

IN THIS SECTION

- [Syntax | 1413](#)
- [Description | 1414](#)
- [Options | 1414](#)
- [Required Privilege Level | 1414](#)
- [Output Fields | 1414](#)
- [Sample Output | 1414](#)
- [Release Information | 1414](#)

Syntax

```
request paragon repair-node address faulty-node-address
```

Description

Repair a faulty node.

Options

address <i>faulty-node-address</i>	IPv4 address of the faulty node

Required Privilege Level

configure

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request paragon repair-node address *faulty-node-address*

```
user@node# request paragon replace-node address 10.1.2.5
Process running with PID: 23xx022
To track progress, run 'monitor start /epic/config/log'
```

Release Information

Command introduced in Juniper Paragon Automation Release 2.0.0.

request paragon replace-node

IN THIS SECTION

● [Syntax](#) | 1415

- [Description | 1415](#)
- [Options | 1415](#)
- [Required Privilege Level | 1415](#)
- [Output Fields | 1415](#)
- [Sample Output | 1416](#)
- [Release Information | 1416](#)

Syntax

```
request paragon replace-node address replacement-node-address
```

Description

Replace a faulty node in your Routing Director cluster with a replacement node. The replacement node IP address can be the same as that of the faulty node or a new one.

Options

address <i>replacement-node-address</i>	IPv4 addresses of the replacement node.

Required Privilege Level

configure

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request paragon replace-node address *replacement-node-address*

```
user@node# request paragon replace-node address 10.1.2.5
Process running with PID: 23xx022
To track progress, run 'monitor start /epic/config/log'
```

Release Information

Command introduced in Juniper Paragon Automation Release 2.0.0.

request paragon restore

IN THIS SECTION

- [Syntax | 1416](#)
- [Description | 1417](#)
- [Options | 1417](#)
- [Required Privilege Level | 1417](#)
- [Output Fields | 1417](#)
- [Sample Output | 1417](#)
- [Release Information | 1418](#)

Syntax

```
request paragon restore backup-id backup-id
```

```
request paragon restore sync backup-id backup-id
```

Description

Restore your Routing Director configuration to an earlier backed-up version.

Options

<code>backup-id <i>backup-id</i></code>	Restore your configuration to the version saved in the backed-up file with the specified <i>backup-id</i> . To view a list of available backup files, see " show paragon backup " on page 1457.
<code>sync backup-id <i>backupID</i></code>	Request Routing Director to reparse the restored configuration.

Required Privilege Level

configure

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request paragon restore backup-id *backup-id*

```
root@node> request paragon restore start backup-id 20240507-082612
Postgres restore Started
Postgres restore completed
Using arango CRDN pod foghorn-crdn-ebopribl-413778
Using arango CRDN pod foghorn-crdn-ebopribl-413778
ArangoDB restore completed
Airflow restore completed
Timescale restore Started
Timescale restore completed
Successfully flushed redis cache
```

request paragon restore sync backup-id *backupID*

```
user@node> request paragon restore sync backup-id 202505103-082640
Insights config reparsing started
Checking Readiness of Config-server
Calling Reparse config for Insights
Reparse config in Insights request is successfull
Insights config reparsing Completed
```

Release Information

Command introduced in Juniper Paragon Automation Release 2.0.0.

sync command modified in Juniper Paragon Automation Release 2.4.0.

RELATED DOCUMENTATION

request paragon backup 1381
show paragon backup 1457
request paragon deploy cluster 1403
request paragon service destroy 1425

request paragon reboot

IN THIS SECTION

- [Syntax | 1419](#)
- [Description | 1419](#)
- [Options | 1419](#)
- [Required Privilege Level | 1419](#)
- [Output Fields | 1419](#)
- [Sample Output | 1419](#)
- [Release Information | 1420](#)

Syntax

```
request paragon reboot type (cluster|node)
```

Description

Reboot a Routing Director node VM or the entire cluster.

Options

type cluster	Reboot all the nodes in the Routing Director cluster .
type node	Reboot the node you are logged in to in the Routing Director cluster.

Required Privilege Level

configure

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system reboot
type node

```
user@node1> request paragon reboot
type node
node1
Rebooting current node and connection will be lost: node1
Connection to node1 closed by remote host.
Connection to node1 closed.
```

request system reboot type cluster

```
user@node1> request paragon reboot
type cluster
node1
Rebooting the node: node2
Finished rebooting the node: node2
Rebooting the node: node3
Finished rebooting the node: node3
Rebooting the node: node4
Finished rebooting the node: node4
Rebooting current node now and connection will be lost: node1
Connection to node1 closed by remote host.
Connection to node1 closed.
```

Release Information

Command introduced in Juniper Paragon Automation Release 2.3.0.

RELATED DOCUMENTATION

| [request paragon shutdown](#) | [1422](#)

request paragon running-config

IN THIS SECTION

- [Syntax](#) | [1421](#)
- [Description](#) | [1421](#)
- [Options](#) | [1421](#)
- [Output Fields](#) | [1421](#)
- [Required Privilege Level](#) | [1421](#)
- [request paragon running-config](#) | [1421](#)

Syntax

```
request paragon running-config
```

Description

In case the **config.yml** and **inventory** files are missing, corrupted, or inadvertently tampered with on any node, generate the Routing Director cluster current configuration files on the node.

Options

This command has no options.

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Required Privilege Level

configure

request paragon running-config

```
user@node> request paragon running-config
Paragon inventory file saved at /epic/config/inventory
Paragon config file saved at /epic/config/config.yml
```

Release Information

Command introduced in Juniper Paragon Automation Release 2.0.0.

request paragon shutdown

IN THIS SECTION

- [Syntax | 1422](#)
- [Description | 1422](#)
- [Options | 1422](#)
- [Required Privilege Level | 1422](#)
- [Output Fields | 1423](#)
- [Sample Output | 1423](#)
- [Release Information | 1423](#)

Syntax

```
request paragon shutdown type (cluster|node)
```

Description

Shut down a Routing Director node VM or the entire cluster.

Options

type cluster	Shut down all the nodes in the Routing Director cluster .
type node	Shut down the node you are logged in to in the Routing Director cluster.

Required Privilege Level

configure

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system shutdown

type node

```
user@node1> request paragon shutdown
type node
node1
Shutting down current node and connection will be lost: node1
Connection to node1 closed by remote host.
Connection to node1 closed.
```

request system shutdown

type cluster

```
user@node2> request paragon shutdown
type cluster
node2
Shutting down the node: node1
The node node1 has been shutdown
Shutting down the node: node3
The node node3 has been shutdown
Shutting down the node: node4
The node node4 has been shutdown
Shutting down current node now and connection will be lost: node2
Connection to node2 closed by remote host.
Connection to node2 closed.
```

Release Information

Command introduced in Juniper Paragon Automation Release 2.3.0.

RELATED DOCUMENTATION

| [request paragon reboot](#) | 1418

request paragon service

IN THIS SECTION

- [Syntax](#) | 1424
- [Description](#) | 1424
- [Options](#) | 1424
- [Required Privilege Level](#) | 1425
- [Output Fields](#) | 1425
- [Sample Output](#) | 1425
- [Release Information](#) | 1425

Syntax

```
request paragon service destroy
```

```
request paragon service start
```

Description

Destroy the Routing Director cluster services before a restore operation. Also, restart all services when needed.

Options

service destroy	Destroy the Routing Director cluster applications and services. This command destroys all services, and might remove user data from the system.
service start	Restart all services.

Required Privilege Level

configure

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request paragon service destroy

```
user@node> request paragon service destroy
Warning! This will destroy all services, which might remove user data from the system!
Do you want to proceed? (Y/N) Y
Proceed with paragon service destroy.
```

Release Information

Command introduced in Juniper Paragon Automation Release 2.4.0.

RELATED DOCUMENTATION

[request paragon restore](#) | 1416

[request paragon deploy cluster](#) | 1403

request paragon ssh

IN THIS SECTION

- [Syntax | 1426](#)
- [Description | 1426](#)
- [Options | 1426](#)
- [Required Privilege Level | 1427](#)
- [Output Fields | 1427](#)
- [Sample Output | 1427](#)
- [Release Information | 1427](#)

Syntax

```
request paragon ssh hostname node-hostname commands-to-execute commands
```

Description

Use the SSH program to open a connection between the local node and a remote node and execute commands on the remote node.

Options

commands-to-execute <i>commands</i>	Enter commands to be executed. Include one or more CLI commands by enclosing them in quotation marks and separating the commands with semicolons. For example, 'cli-command1 ; cli-command2'. Examples of commands that can be executed include, df -h, ls, echo, cp, kubectl get configmaps -n <i>namespace</i> , kubectl get pods -A, and so on.
hostname <i>node_hostname</i>	Hostname or address of the remote node.

Required Privilege Level

configure

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request paragon ssh hostname *node-hostname* commands-to-execute *commands*

```
user@node> request paragon ssh hostname node2 commands-to-execute ls
controller.sh
epic
prep.sh
set_hostname.sh
setup_key.py
snap
sources.list
start.sh
temp
troubleshooting
utils.sh

user@node>
```

Release Information

Command introduced in Juniper Paragon Automation Release 2.0.0.

request paragon ssh-key

IN THIS SECTION

● [Syntax](#) | 1428

- [Description | 1428](#)
- [Options | 1428](#)
- [Output Fields | 1428](#)
- [Required Privilege Level | 1429](#)
- [Sample Output | 1429](#)
- [Release Information | 1430](#)

Syntax

```
request paragon ssh-key
```

```
request paragon ssh-key username username password password nodes IP-1, IP-2, IP-3...
```

Description

Generate the SSH key on all the Routing Director cluster nodes and add public keys.

Options

none	Generate the SSH key for the Routing Director cluster nodes.
nodes <i>IP-1, IP-2, IP-3...</i>	Comma separated list of IPv4 addresses of the nodes in the cluster.
password <i>password</i>	Password to log in to the nodes.
username <i>username</i>	Username to log in to the nodes. For example, root.

Output Fields

When you enter this command, you are provided feedback on the status and output of your request.

Required Privilege Level

configure

Sample Output

request paragon ssh-key

```
root@eop> request paragon ssh-key
Please enter comma-separated list of IP addresses: 10.1.2.3,10.1.2.4,10.1.2.5,10.1.2.6
Please enter SSH username for the node(s): root
Please enter SSH password for the node(s): rootpassword
checking server reachability and ssh connectivity ...
Connectivity ok for 10.1.2.3
Connectivity ok for 10.1.2.4
Connectivity ok for 10.1.2.5
Connectivity ok for 10.1.2.6
SSH key pair already exists in 10.1.2.3
SSH key pair already exists in 10.1.2.4
SSH key pair already exists in 10.1.2.5
SSH key pair already exists in 10.1.2.6
copied from 10.1.2.3 to 10.1.2.3
copied from 10.1.2.3 to 10.1.2.4
copied from 10.1.2.3 to 10.1.2.5
copied from 10.1.2.3 to 10.1.2.6
copied from 10.1.2.4 to 10.1.2.3
copied from 10.1.2.4 to 10.1.2.4
copied from 10.1.2.4 to 10.1.2.5
copied from 10.1.2.4 to 10.1.2.6
copied from 10.1.2.5 to 10.1.2.3
copied from 10.1.2.5 to 10.1.2.4
copied from 10.1.2.5 to 10.1.2.5
copied from 10.1.2.5 to 10.1.2.6
copied from 10.1.2.6 to 10.1.2.3
copied from 10.1.2.6 to 10.1.2.4
copied from 10.1.2.6 to 10.1.2.5
copied from 10.1.2.6 to 10.1.2.6
```

request paragon ssh-key username *username* password *password* nodes *IP-1, IP-2, IP-3..*

```

root@eop> request paragon ssh-key username root password rootpassword nodes
10.1.2.3,10.1.2.4,10.1.2.5,10.1.2.6
checking server reachability and ssh connectivity ...
Connectivity ok for 10.1.2.3
Connectivity ok for 10.1.2.4
Connectivity ok for 10.1.2.5
Connectivity ok for 10.1.2.6
SSH key pair already exists in 10.1.2.3
SSH key pair already exists in 10.1.2.4
SSH key pair already exists in 10.1.2.5
SSH key pair already exists in 10.1.2.6
copied from 10.1.2.3 to 10.1.2.3
copied from 10.1.2.3 to 10.1.2.4
copied from 10.1.2.3 to 10.1.2.5
copied from 10.1.2.3 to 10.1.2.6
copied from 10.1.2.4 to 10.1.2.3
copied from 10.1.2.4 to 10.1.2.4
copied from 10.1.2.4 to 10.1.2.5
copied from 10.1.2.4 to 10.1.2.6
copied from 10.1.2.5 to 10.1.2.3
copied from 10.1.2.5 to 10.1.2.4
copied from 10.1.2.5 to 10.1.2.5
copied from 10.1.2.5 to 10.1.2.6
copied from 10.1.2.6 to 10.1.2.3
copied from 10.1.2.6 to 10.1.2.4
copied from 10.1.2.6 to 10.1.2.5
copied from 10.1.2.6 to 10.1.2.6

```

Release Information

Command introduced in Juniper Paragon Automation Release 2.0.0.

request paragon storage cleanup

IN THIS SECTION

- [Syntax | 1431](#)
- [Description | 1431](#)
- [Options | 1431](#)
- [Required Privilege Level | 1432](#)
- [Sample Output | 1432](#)
- [Release Information | 1433](#)

Syntax

```
request paragon storage cleanup
```

Description

Deletes all core files and sub directories, local_host core files, logs, traces, and all files under **/var/log** and **/var/crash** directories.

Options

none	Clears up storage files.
dry-run	Lists all the files to be cleared up on all cluster nodes but does not remove them.
extra-path <i>path</i>	Extra path on the nodes (if the path exists), separated by a space and enclosed in quotes. For example: "/foo /bar/zot" .

Required Privilege Level

configure

Sample Output

request paragon storage cleanup dry-run

```

user@node> request paragon storage cleanup dry-run
warning: *** operating on 10.1.2.3 ***
contents to clean up
Node: node
List of all core files to be cleared:
      Size      Date      Name
List of local_host core files to be cleared:
      Size      Date      Name
List of core sub directory files to be cleared:
      Size      Date      Name
Clears all logs, traces, /var/log/*, and /var/crash/*
      Size      Date      Name
60K Fri Apr 12 13:58 /var/log/action.log
8.0K Wed Apr  3 22:27 /var/log/alternatives.log
0 Wed Apr  3 21:58 /var/log/btmp
16K Thu Apr 11 10:23 /var/log/cscript.log
12K Wed Apr  3 22:27 /var/log/dpkg.log
0 Wed Apr  3 22:27 /var/log/license
0 Tue Apr  9 09:54 /var/log/messages
0 Tue Apr  9 09:54 /var/log/security
4.0K Fri Apr 12 13:06 /var/log/wtmp
-rw-r--r-- 1 root root 404226909 Apr  4 04:05 /root/epic/3rdparty.tar.gz
-rw-r--r-- 1 root root 455199761 Apr  4 04:05 /root/epic/davinci.tar.gz

/root/epic/temp/:
total 0
dry-run was enabled, nothing was removed
dry run docker system prune
warning: *** operation on 10.1.2.3 succeeds ***
warning: *** operating on 10.1.2.4 ***
contents to clean up
Node: node2
List of all core files to be cleared:

```

```

      Size      Date      Name
List of local_host core files to be cleared:
      Size      Date      Name
List of core sub directory files to be cleared:
      Size      Date      Name
Clears all logs, traces, /var/log/*, and /var/crash/*
      Size      Date      Name
8.0K Tue Apr  9 10:38 /var/log/action.log
8.0K Wed Apr  3 22:27 /var/log/alternatives.log
0 Wed Apr  3 21:58 /var/log/btmp
12K Thu Apr 11 10:23 /var/log/cscript.log
12K Wed Apr  3 22:27 /var/log/dpkg.log
0 Wed Apr  3 22:27 /var/log/license
0 Thu Apr 11 10:23 /var/log/messages
0 Thu Apr 11 10:23 /var/log/security
0 Wed Apr  3 21:58 /var/log/wtmp
-rw-r--r-- 1 root root 404226909 Apr  4 04:05 /root/epic/3rdparty.tar.gz
-rw-r--r-- 1 root root 455199761 Apr  4 04:05 /root/epic/davinci.tar.gz

<output snipped>

```

Release Information

Command introduced in Juniper Paragon Automation Release 2.0.0.

request paragon super-user password reset

IN THIS SECTION

- [Syntax | 1434](#)
- [Description | 1434](#)
- [Options | 1434](#)
- [Output Fields | 1434](#)
- [Required Privilege Level | 1434](#)
- [Release Information | 1434](#)

Syntax

```
request paragon super-user password reset user superuser_email_address password
new_temporary_password
```

Description

Reset the password of a superuser of the Web GUI using Paragon Shell when SMTP is not configured. If a super user of the GUI forgets their password, a Paragon Shell superuser (typically a system administrator) can reset the password using this command. The Paragon Shell superuser shares the temporary password with the GUI superuser who can use the temporary password to log in to the GUI and change it as required.

Options

password <i>new_temporary_password</i>	The temporary password.
user <i>superuser_email_address</i>	E-mail address of the superuser whose password will be reset.

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Required Privilege Level

configure

Release Information

Command introduced in Juniper Paragon Automation Release 2.0.0.

request system decrypt password

IN THIS SECTION

- [Syntax | 1435](#)
- [Description | 1435](#)
- [Options | 1435](#)
- [Required Privilege Level | 1435](#)
- [Output Fields | 1435](#)
- [Sample Output | 1436](#)
- [Release Information | 1436](#)

Syntax

```
request system decrypt password
```

Description

Use to display plain text versions of obfuscated (\$9) or encrypted (\$8) passwords. If the password was encrypted using the new \$8\$ method, you are prompted for the primary password.

Options

<i>encrypted_password</i>	Decrypt a \$8\$-encrypted or \$9\$-encrypted password.

Required Privilege Level

configure

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system decrypt password *encrypted_password*

```
user@node> request system decrypt password $9$ABC123
Plaintext password: mysecret
```

request system decrypt password *encrypted_password*

```
user@host> request system decrypt password $8$ABC123
Master password:
Plaintext password: mysecret
(Simple passwords like "mysecret" are discouraged. This is an example only.)
```

Release Information

Command introduced in Juniper Paragon Automation Release 2.0.0.

request system reboot

IN THIS SECTION

- [Syntax | 1437](#)
- [Description | 1437](#)
- [Options | 1437](#)
- [Required Privilege Level | 1437](#)
- [Output Fields | 1437](#)
- [Sample Output | 1437](#)
- [Release Information | 1437](#)

Syntax

```
request system reboot
```

Description

Reboot the Routing Director node VM.

Options

This command has no options.

Required Privilege Level

configure

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system reboot

```
user@node> request system reboot

Reboot the system ? [yes,no] (no) yes
user@node> Connection to node closed by remote host.

Connection to node closed.
```

Release Information

Command introduced in Juniper Paragon Automation Release 2.0.0.

request paragon trust data update hostname

IN THIS SECTION

- Syntax | 1438
- Description | 1438
- Options | 1438
- Sample Output | 1438
- Release Information | 1439

Syntax

```
request paragon trust data update hostname hostname
```

Description

Update your static Trust data in an air-gapped environment where your on-premise Routing Director installation does not have access to the Internet. For more information, see ["Update Trust KPIs Using the Paragon Shell CLI Command"](#) on page 430.

Options

<i>hostname</i>	Hostname of the controller cluster node where you update the Trust data.

Sample Output

```
root@node> request paragon trust data update hostname hostname

HPA for trust, network minReplicas is set to: 1 and maxReplicas is set to 1.
horizontalpodautoscaler.autoscaling/network patched Replica count for trust, network is set to:
0 in deployment. deployment.apps/network scaled Deployment for trust, network runmode is set to:
normal. deployment.apps/network configured
```

Persistent Volume/Claim File creation

Persistent Volume/Claim File creation

Persistent Volume/Claim creation. persistentvolume/network-pv created

Persistent Volume/Claim creation. persistentvolumeclaim/network-pvc created

Deployment for trust, network has been patched to include node affinity. deployment.apps/network patched Replica count for trust, network is set to: 1 in deployment. deployment.apps/network scaled Update should now be complete, cleaning up kubernetes settings and scaling back up.

<Output Snipped>

Trust static data update complete

Release Information

Command introduced in Juniper Routing Director Release 2.5.0.

show configuration paragon cluster

IN THIS SECTION

- [Syntax | 1440](#)
- [Description | 1440](#)
- [Options | 1440](#)
- [Required Privilege Level | 1446](#)
- [Output Fields | 1446](#)
- [Sample Output | 1446](#)
- [Release Information | 1447](#)

Syntax

```
show configuration paragon cluster
```

```
show configuration paragon cluster applications (none|applications_options)
```

```
show configuration paragon cluster common-services (none|common_services_options)
```

```
show configuration paragon cluster insights victoria-metrics (none|victoria_metrics_options)
```

```
show configuration paragon cluster install (none|offline-install|prefer-ipv6)
```

```
show configuration paragon cluster kubernetes (none|address-family [none|cluster-ipv6-enabled])
```

```
show configuration paragon cluster mail-server (none|mail_server_options)
```

```
show configuration paragon cluster nodes (none|nodes_options)
```

```
show configuration paragon cluster ntp (none|ntp-servers)
```

```
show configuration paragon cluster papi (none|papi-local-user-management)
```

Description

Display all the configured parameters under the paragon hierarchy.

Options

none	Show all the configured parameters under the paragon hierarchy.

applications (none <i>applications_options</i>)	Show Paragon Active Assurance Test Agent gateway virtual IP address, hostname, and default Web administrator login credentials. <i>applications_options</i> options include: • active-assurance • active-assurance test-agent-gateway-vip • active-assurance test-agent-gateway-vip-ipv6 • active-assurance test-agent-gateway-hostname • aiops • aiops enable-device-health • aiops install-aiml • gnmi-term • gnmi-term gnmi-term-host • gnmi-term gnmi-term-host-ipv6 • oc-term • oc-term oc-term-host • oc-term oc-term-host-ipv6 • oc-term oc-term-previous-management-ip-family • pathfinder pce-server pce-server-vip • pathfinder pce-server pce-server-global-default-tls-mode • pathfinder pce-server pce-server-user-certificate pce-server-use-custom-certificate • pathfinder pce-server pce-server-user-certificate pce-server-custom-certificate-filename • pathfinder pce-server pce-server-user-certificate pce-server-custom-certificate-key-filename

- routingbot
- routingbot install-routingbot
- routingbot routingbot-crpd-vip
- web-ui
- web-ui web-admin-password
- web-ui web-admin-user

For more information, see "[set paragon cluster applications](#)" on page 1498 and "[show configuration paragon cluster applications \(o\)](#)" on page 1448.

common-services (none <i>common_services_options</i>)	Show the common ingress virtual IP (VIP) address, hostname, and custom user certificates usage. <i>common_services_options</i> options include: • ingress ingress-hostname • ingress ingress-vip • ingress ingress-vip-ipv6 • ingress user-certificate • ingress user-certificate use-user-certificate • ingress user-certificate user-certificate-filename • ingress user-certificate-key-filename • metallb metallb-bgp-peer peer-ip <i>v4-address</i> • metallb metallb-bgp-peer peer-ip <i>v4-address</i> peer-asn • metallb metallb-bgp-peer peer-ip <i>v4-address</i> local-asn • metallb metallb-bgp-peer peer-ip <i>v4-address</i> local-nodes • metallb metallb-bgp-peer peer-ip <i>v6-address</i> • metallb metallb-bgp-peer peer-ip <i>v6-address</i> peer-asn • metallb metallb-bgp-peer peer-ip <i>v6-address</i> local-asn • metallb metallb-bgp-peer peer-ip <i>v6-address</i> local-nodes Where <i>v4-address</i> and <i>v6-address</i> is the IP and IPv6 address of the top-of-rack (TOR) router respectively. For more information, see "set paragon cluster common-services" on page 1504 and "show configuration paragon cluster common-services (o)" on page 1451.

insights (none victoria-metrics <i>victoria_metrics_options</i>)	Display Victoria Metrics configuration parameters. <i>victoria_metrics_options</i> options include: • vm-cluster-replication-factor • vm-cluster-retention-period • vm-cluster-storage-instances For more information, see "set paragon cluster insights victoria-metrics" on page 1507.
install (none offline-install prefer-ipv6)	Display the cluster installation parameters and whether the Routing Director cluster is enabled to be installed in an air-gap environment. For more information, see "set paragon cluster install" on page 1509.
kubernetes (none address-family [none cluster-ipv6-enabled])	Display whether the usage of IPv6 addresses is enabled or disabled for the cluster. For more information, see "set paragon cluster kubernetes" on page 1511.
mail-server (none <i>mail_server_options</i>)	Display SMTP configuration settings. <i>mail_server_options</i> options include: • smtp-allowed-sender-domains • smtp-relayhost • smtp-relayhost-password • smtp-relayhost-username • smtp-sender-email • smtp-sender-name For more information, see "set paragon cluster mail-server" on page 1512.

nodes (none <i>nodes_options</i>)	Display IP address or hostname of the Routing Director cluster node and the index number of the node. <i>nodes_options</i> options include: • <i>kubernetes index</i> • <i>kubernetes index</i> address • <i>kubernetes index</i> hostname For more information, see "set paragon cluster nodes" on page 1515.
ntp (none ntp-servers)	Display the configured NTP servers. For more information, see "set paragon cluster ntp" on page 1516.
papi (none papi-local-user-management)	Display whether local user management on the nodes is enabled or disabled. For more information, see "set paragon cluster papi" on page 1518.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided with information about your Routing Director cluster.

Sample Output

show configuration paragon cluster

```

root@node> show configuration paragon cluster
nodes {
  kubernetes 1 {
    address 10.1.2.3;
  }
}
```

```

    kubernetes 2 {
        address 10.1.2.4;
    }
    kubernetes 3 {
        address 10.1.2.5;
    }
    kubernetes 4 {
        address 10.1.2.6;
    }
}
ntp {
    ntp-servers pool.ntp.org;
}
common-services {
    ingress {
        ingress-vip 10.1.2.7;
    }
}
applications {
    active-assurance {
        test-agent-gateway-vip 10.1.2.8;
    }
    web-ui {
        web-admin-user "user@juniper.net";
        web-admin-password ENC:U2FsdGVkX201aFjoisYla0aaaaaqLmNvcLumKS1uUoU=;
    }
}
}

```

Release Information

Statement introduced in Juniper Paragon Automation Release 2.0.0.

pathfinder option introduced in Juniper Paragon Automation Release 2.1.0.

insights option introduced in Juniper Paragon Automation Release 2.1.0.

kubernetes option introduced in Juniper Paragon Automation Release 2.2.0.

prefer-ipv6 option introduced in Juniper Paragon Automation Release 2.2.0.

ingress ingress-vip-ipv6 option introduced in Juniper Paragon Automation Release 2.2.0.

test-agent-gateway-vip-ipv6 option introduced in Juniper Paragon Automation Release 2.2.0.

aiops, gnmi-term, oc-term, routingbot, and metallb options introduced in Juniper Paragon Automation Release 2.4.0.

pce-server-global-default-tls-mode and pce-server-user-certificate options introduced in Juniper Routing Director Release 2.5.0.

RELATED DOCUMENTATION

[delete paragon cluster | 1491](#)

[show configuration paragon cluster applications \(o\) | 1448](#)

[show configuration paragon cluster common-services \(o\) | 1451](#)

show configuration paragon cluster applications (o)

IN THIS SECTION

- [Syntax | 1449](#)
- [Description | 1449](#)
- [Options | 1449](#)
- [Required Privilege Level | 1450](#)
- [Output Fields | 1450](#)
- [Release Information | 1450](#)

Syntax

```
show configuration paragon cluster applications
```

```
show configuration paragon cluster applications active-assurance (none|test-agent-gateway-vip|  
test-agent-gateway-vip-ipv6|test-agent-gateway-hostname)
```

```
show configuration paragon cluster applications aiops (none|enable-device-health|install-aiml)
```

```
show configuration paragon cluster applications gnmi-term (none|gnmi-term-host|gnmi-term-host-  
ipv6)
```

```
show configuration paragon cluster applications oc-term (none|oc-term-host|oc-term-host-ipv6|oc-  
term-previous-management-ip-family)
```

```
show configuration paragon cluster applications pathfinder pce-server pce-server-vip
```

```
show configuration paragon cluster applications routingbot (none|install-routingbot|routingbot-  
crpd-vip)
```

```
show configuration paragon cluster applications web-ui (none|web-admin-password|web-admin-user)
```

Description

Show all configured applications details.

Options

none	Show all configured applications details.

active-assurance (none test-agent-gateway-vip test-agent-gateway-vip-ipv6 test-agent-gateway-hostname)	Show all the configured Paragon Active Assurance Test Agent gateway details, such as IPv4 and IPv6 addresses, and hostname.
aiops (none enable-device-health install-aiml)	Show AI/ML (artificial intelligence [AI] and machine learning [ML]) configuration settings including device-health monitoring settings.
gnmi-term (none gnmi-term-host gnmi-term-host-ipv6)	Show the configured IPv4 or IPv6 VIP addresses for gNMI.
oc-term (none oc-term-host oc-term-host-ipv6 oc-term-previous-management-ip-family)	Show the configured IPv4 or IPv6 VIP addresses for OC-TERM.
pathfinder pce-server pce-server-vip	Show the configured PCE server VIP address details.
routingbot (none install-routingbot routingbot-crpd-vip)	Show routing observability feature configuration settings including the CRPD VIP address.
web-ui (none web-admin-password web-admin-user)	Show the Web administrator login e-mail ID and password. The password is encrypted. The Web administrator is the first user to log in to the Routing Director GUI and has super user privileges.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided with information about your Routing Director cluster.

Release Information

Statement introduced in Juniper Paragon Automation Release 2.0.0.

test-agent-gateway-vip-ipv6 option introduced in Juniper Paragon Automation Release 2.2.0.

pathfinder option introduced in Juniper Paragon Automation Release 2.1.0.

aiops, gnmi-term, oc-term, and routingbot options introduced in Juniper Paragon Automation Release 2.4.0.

RELATED DOCUMENTATION

| [show configuration paragon cluster](#) | 1439

show configuration paragon cluster common-services (o)

IN THIS SECTION

- [Syntax](#) | 1452
- [Description](#) | 1452
- [Options](#) | 1452
- [Required Privilege Level](#) | 1454
- [Output Fields](#) | 1454
- [Release Information](#) | 1454

Syntax

```
show configuration paragon cluster common-services
```

```
show configuration paragon cluster common-services ingress (none|ingress-hostname|ingress-vip|ingress-vip-ipv6|user-certificate)
```

```
show configuration paragon cluster common-services ingress user-certificate use-user-certificate (none|true|false)
```

```
show configuration paragon cluster common-services ingress user-certificate user-certificate-filename
```

```
show configuration paragon cluster common-services ingress user-certificate user-certificate-key-filename
```

```
show configuration paragon cluster common-services metallb metallb-bgp-peer peer-ip v4-address (none|peer-asn|local-asn|local-nodes)
```

```
show configuration paragon cluster common-services metallb metallb-bgp-peer-ipv6 peer-ip v6 address (none|peer-asn|local-asn|local-nodes)
```

Description

Show all common ingress configuration information.

Options

(none ingress)	Show all common ingress configuration information.

ingress ingress-hostname	Show the hostname for the ingress VIP address.
ingress ingress-vip	Show the common ingress IPv4 VIP address.
ingress ingress-vip-ipv6	Show the common ingress IPv6 VIP address.
ingress user-certificate	Show custom user certificates configuration.
ingress user-certificate use-user-certificate (none true false)	Show custom user certificates configuration.
ingress user-certificate user-certificate-filename	Show the custom user certificate filename.
ingress user-certificate-key-filename	Show the custom user certificate key filename.
metallb metallb-bgp-peer peer-ip <i>v4-address</i>	Show all the configured BGP peering information between the specified top-of-rack (ToR) router and the cluster nodes connected to it. Here, <i>v4-address</i> is the IP address of the ToR router.
metallb metallb-bgp-peer peer-ip <i>v4-address</i> peer-asn	Show the configured AS number for the specified ToR router.
metallb metallb-bgp-peer peer-ip <i>v4-address</i> local-asn	Show the configured AS number for the cluster nodes connected to the specified ToR router.
metallb metallb-bgp-peer peer-ip <i>v4-address</i> local-nodes	Show the cluster nodes connected to the specified ToR router.
metallb metallb-bgp-peer-ipv6 peer-ip <i>v6 address</i>	Show all the configured BGP peering information between the specified top-of-rack (ToR) router and the cluster nodes connected to it. Here, <i>v6 address</i> is the IPv6 address of the ToR router.
metallb metallb-bgp-peer-ipv6 peer-ip <i>v6 address</i> peer-asn	Show the configured AS number for the specified ToR router.

<code>metallb metallb-bgp-peer-ipv6 peer-ip v6 address local-asn</code>	Show the configured AS number for the cluster nodes connected to the specified ToR router.
<code>metallb metallb-bgp-peer-ipv6 peer-ip v6 address local-nodes</code>	Show the cluster nodes connected to the specified ToR router.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided with information about your Routing Director cluster.

Release Information

Statement introduced in Juniper Paragon Automation Release 2.0.0.

ingress-vip-ipv6 option introduced in Juniper Paragon Automation Release 2.2.0.

metallb option added in Juniper Paragon Automation Release 2.4.0.

RELATED DOCUMENTATION

| [show configuration paragon cluster](#) | 1439

show host disk usage

IN THIS SECTION

- [Syntax](#) | 1455
- [Description](#) | 1455
- [Options](#) | 1455
- [Required Privilege Level](#) | 1455

- [Output Fields | 1455](#)
- [Sample Output | 1455](#)
- [Release Information | 1456](#)

Syntax

```
show host disk usage node address
```

Description

Show disk space usage of the specified node of the Routing Director cluster.

Options

<code>node <i>address</i></code>	IPv4 address of the node for which you want to view disk usage.

Required Privilege Level

configure

Output Fields

When you enter this command, the disk space usage of the node is displayed.

Sample Output

show host disk usage

```
user@node> show host disk usage node 10.1.2.3
Filesystem
Size  Used Avail Use% Mounted on
tmpfs
```



```

3.2G  12M  3.2G  1% /run
/dev/
sda2
148G  66G  75G  47% /
tmpfs
16G    0   16G   0% /dev/shm
tmpfs
5.0M    0  5.0M   0% /run/lock
/dev/mapper/
export
98G   23G   71G  24% /export
tmpfs
3.2G  8.0K  3.2G   1% /run/user/0
overlay
148G  66G   75G  47% /var/lib/docker/
overlay2/26aab2eadcb885446a4bc4f5aeddd14770241a9206fc405f877632b4f79f13ae/merged
overlay
148G  66G   75G  47% /var/lib/docker/
overlay2/7fed029594e553f25fd7b3e92ec1e85ea67480c0ba2dbf5ba71ab427a29ae524/merged
shm
64M    0   64M   0% /run/k3s/containerd/
io.containerd.grpc.v1.cri/sandboxes/
f133003ba97c483e278768ee8eaa5b1818f1f9563b920f745574231dd6cd88b4/shm
overlay
148G  66G   75G  47% /run/k3s/containerd/
io.containerd.runtime.v2.task/k8s.io/
f133003ba97c483e278768ee8eaa5b1818f1f9563b920f745574231dd6cd88b4/rootfs
overlay
148G  66G   75G  47% /run/k3s/containerd/
io.containerd.runtime.v2.task/k8s.io/
2795f4f76c5e0ff2caa11ea57ee93325ac9a54259432598cf3b47eb83ca671e0/rootfs
shm
<output snipped>

```

Release Information

Command introduced in Juniper Paragon Automation Release 2.0.0.

show paragon backup

IN THIS SECTION

- [Syntax | 1457](#)
- [Description | 1458](#)
- [Options | 1458](#)
- [Required Privilege Level | 1458](#)
- [Output Fields | 1458](#)
- [Sample Output | 1459](#)
- [Release Information | 1461](#)

Syntax

```
show paragon backup
```

```
show paragon backup remote storage-location remote-path user username password password
```

```
show paragon backup download status backup-id backup-ID
```

```
show paragon backup include-failure (true | false)
```

```
show paragon backup status backup-id backup-ID
```

```
show paragon backup upload status backup-id backup-ID
```

```
show paragon backup schedule
```

Description

Show the backup folders available across all the nodes of the Routing Director cluster, and also the backup folders available in a remote location. You can also view the status of the backup operation and the upload and download status of the backup folders from the remote location. Additionally, you can view any configured backup schedule.

Options

none	Display all backup folders available across all nodes of the Routing Director cluster.
remote storage-location <i>remote-path</i> user <i>username</i> password <i>password</i>	Display all backup folders available in the specified by the remote-path in a remote server location. <i>username</i> and <i>password</i> are the login credentials to log in to the remote server.
download status backup-id <i>backup-ID</i>	When downloading a backed-up folder from a remote location, display the download status of the folder.
include-failure (true false)	Display the failed backup attempts across all nodes when set to <i>true</i> .
status backup-id <i>backup-ID</i>	Display the latest backup command status.
upload status backup-id <i>backup-ID</i>	When uploading a backed-up folder to a remote location, display the upload status of the folder.
schedule	View any configured backup schedule.

Required Privilege Level

configure

Output Fields

[Table 341 on page 1459](#) lists the output fields for the `show paragon backup` command.

Table 341: show paragon backup Output Fields

Field Name	Field Description
Node	IP address of the node on which the backup folder is located.
Backup ID	Unique identifier for the backup folder name.

Sample Output

show paragon backup

```

user@node> show paragon backup
Establishing SSH connection with 10.1.2.5
Establishing SSH connection with 10.1.2.6
Establishing SSH connection with 10.1.2.7

Following successful backups are available:
      Node      Backup ID
      10.1.2.4      20240507-082612
      10.1.2.5      20240507-083155

```

show paragon backup include-failure false

```

user@node> show paragon backup include-failure false
Establishing SSH connection with 10.1.2.5
Establishing SSH connection with 10.1.2.6
Establishing SSH connection with 10.1.2.7

Following successful backups are available:
      Node      Backup ID
      10.1.2.4      20240507-082612
      10.1.2.5      20240507-083155

```

show paragon backup include-failure true

```
user@node> show paragon backup include-failure true
```

```
Establishing SSH connection with 10.1.2.5
```

```
Establishing SSH connection with 10.1.2.6
```

```
Establishing SSH connection with 10.1.2.7
```

```
Following successful backups are available:
```

Node	Backup ID
10.1.2.4	20240507-082612
10.1.2.5	20240507-083155

```
Establishing SSH connection with 10.1.2.5
```

```
Establishing SSH connection with 10.1.2.6
```

```
Establishing SSH connection with 10.1.2.7
```

```
Following un-successful backups are available:
```

Node	Backup ID
10.1.2.6	20240507-083613

show paragon backup status backup-id 20241111-050227

```
user@node> show paragon backup status backup-id 20241111-050227
```

```
Establishing SSH connection with 10.1.2.5
```

```
Establishing SSH connection with 10.1.2.6
```

```
Establishing SSH connection with 10.1.2.7
```

```
Backup status for backup ID 20241111-050227:
```

```
Paragon configuration : SUCCESS
```

```
Timescale job status : SUCCESS
```

```
Opensearch job status : SUCCESS
```

```
Victoriametrics job status : SUCCESS
```

show paragon backup download status backup-id 20241111-050227

```
user@node> show paragon backup download status backup-id 20241111-050227
```

```
Download job status for backup ID 20241111-050227:
```

```
Configuration data download status : SUCCESS
```

```
Opensearch data download status : SUCCESS
Victoriametrics data download status : SUCCESS
```

show paragon backup upload status backup-id 20241111-050227

```
user@node> show paragon backup upload status backup-id 20241111-050227
Upload job status for backup ID 20241111-050227:
  Configuration data upload status : SUCCESS
  Opensearch data upload status : SUCCESS
  Victoriametrics data upload status : SUCCESS
```

show paragon backup schedule

```
root@node> show paragon backup schedule
Backup scheduled on node VM1 : At 12:00 AM, on day 3 of the month
Backup type: telemetry
Remote storage location: scp://10.1.2.3/tmp/backup
Remote user: root
```

Release Information

Command introduced in Juniper Paragon Automation Release 2.0.0.

download, upload, and status options introduced in Juniper Paragon Automation Release 2.3.0.

remote option introduced in Juniper Paragon Automation Release 2.4.0.

schedule option introduced in Juniper Routing Director Release 2.5.0.

RELATED DOCUMENTATION

[request paragon backup](#) | **1381**

[request paragon restore](#) | **1416**

show paragon certificate expiry-date certificate-type

IN THIS SECTION

- [Syntax | 1462](#)
- [Description | 1462](#)
- [Options | 1462](#)
- [Required Privilege Level | 1463](#)
- [Output Fields | 1463](#)
- [Sample Output | 1463](#)
- [Release Information | 1464](#)

Syntax

```
show paragon certificate expiry-date certificate-type (ingress-cert|kubernetes-cert)
```

Description

Display the expiry dates for the Kubernetes ingress and RKE2 certificates.

Options

Table 342:

ingress-cert	Displays the Kubernetes ingress certificate and, if available, the custom user ingress certificate.
kubernetes-cert	Displays whether the cluster node is a server or agent node and displays the respective Kubernetes certificates.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status and output of your request.

Sample Output

show paragon certificate expiry-date certificate-type ingress-cert

```
user@node> show paragon certificate expiry-date certificate-type ingress-cert
The kubernetes ingress certificates are as follows:
  cfssl-common-cert      2034-04-06T18:29:25Z
  ingress-default-cert   2024-07-07T18:12:14Z
  ingress-default-cert   2024-07-07T18:11:51Z
Secret 'ingress-user-cert' does not exist in namespace 'traefik'. Hence, there are no customer
owned certificates!
```

show paragon certificate expiry-date certificate-type kubernetes-cert

```
user@node> show paragon certificate expiry-date certificate-type kubernetes-cert
The host_name where certificates are checked is: node
We have both server and agent paths available in the host
The certificates list is:
  /var/lib/rancher/rke2/server/tls/client-admin.crt
notAfter=May  2 06:57:10 2025 GMT
  /var/lib/rancher/rke2/server/tls/client-auth-proxy.crt
notAfter=May  2 06:57:10 2025 GMT
  /var/lib/rancher/rke2/server/tls/client-ca.crt
notAfter=Apr 30 06:57:10 2034 GMT
  /var/lib/rancher/rke2/server/tls/client-ca.nochain.crt
notAfter=Apr 30 06:57:10 2034 GMT
  /var/lib/rancher/rke2/server/tls/client-controller.crt
notAfter=May  2 06:57:10 2025 GMT
  /var/lib/rancher/rke2/server/tls/client-kube-apiserver.crt
notAfter=May  2 06:57:10 2025 GMT
  /var/lib/rancher/rke2/server/tls/client-kube-proxy.crt
notAfter=May  2 06:57:10 2025 GMT
```



```

/var/lib/rancher/rke2/server/tls/client-rke2-cloud-controller.crt
notAfter=May  2 06:57:10 2025 GMT
/var/lib/rancher/rke2/server/tls/client-rke2-controller.crt
notAfter=May  2 06:57:10 2025 GMT
/var/lib/rancher/rke2/server/tls/client-scheduler.crt
notAfter=May  2 06:57:10 2025 GMT
/var/lib/rancher/rke2/server/tls/client-supervisor.crt
notAfter=May  2 06:57:10 2025 GMT
/var/lib/rancher/rke2/server/tls/request-header-ca.crt
notAfter=Apr 30 06:57:10 2034 GMT
/var/lib/rancher/rke2/server/tls/server-ca.crt
notAfter=Apr 30 06:57:10 2034 GMT
/var/lib/rancher/rke2/server/tls/server-ca.nochain.crt
notAfter=Apr 30 06:57:10 2034 GMT
/var/lib/rancher/rke2/server/tls/serving-kube-apiserver.crt
notAfter=May  2 06:57:10 2025 GMT
/var/lib/rancher/rke2/agent/client-ca.crt
notAfter=Apr 30 06:57:10 2034 GMT
/var/lib/rancher/rke2/agent/client-kubelet.crt
notAfter=May  2 06:57:11 2025 GMT
/var/lib/rancher/rke2/agent/client-kube-proxy.crt
notAfter=May  2 06:57:10 2025 GMT
/var/lib/rancher/rke2/agent/client-rke2-controller.crt
notAfter=May  2 06:57:10 2025 GMT
/var/lib/rancher/rke2/agent/server-ca.crt
notAfter=Apr 30 06:57:10 2034 GMT
/var/lib/rancher/rke2/agent/serving-kubelet.crt
notAfter=May  2 06:57:11 2025 GMT

```

Release Information

Command introduced in Juniper Paragon Automation Release 2.0.0.

show paragon cluster (o)

IN THIS SECTION

- [Syntax | 1465](#)
- [Description | 1465](#)
- [Options | 1466](#)
- [Required Privilege Level | 1466](#)
- [Output Fields | 1466](#)
- [Sample Output | 1466](#)
- [Release Information | 1466](#)

Syntax

```
show paragon cluster
```

```
show paragon cluster pods
```

```
show paragon cluster nodes
```

```
show paragon cluster namespaces
```

```
show paragon cluster details
```

Description

Show information of your Routing Director cluster.

Options

none	Display cluster information
pods	Display pod information of your Routing Director cluster.
nodes	Display node information of your Routing Director cluster.
namespaces	Display namespace information of your Routing Director cluster.
details	Display storage and controller node information of your Routing Director cluster.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided with information about your Routing Director cluster.

Sample Output

show paragon cluster

```
user@node> show paragon cluster
Kubernetes master is running at https://10.1.2.3:6443
CoreDNS is running at https://10.1.2.3:6443/api/v1/namespaces/kube-system/services/rke2-coredns-
rke2-coredns:udp-53/proxy
```

Release Information

Command introduced in Juniper Paragon Automation Release 2.0.0.

show paragon cluster details

IN THIS SECTION

- [Syntax | 1467](#)
- [Description | 1467](#)
- [Options | 1467](#)
- [Required Privilege Level | 1467](#)
- [Output Fields | 1467](#)
- [Sample Output | 1468](#)
- [Release Information | 1468](#)

Syntax

```
show paragon cluster details
```

Description

Show the storage and controller nodes of your Routing Director cluster.

Options

This command has no options.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided with a list of IP addresses of the controller and storage nodes of your Routing Director cluster.

Sample Output

show paragon cluster details

```
user@node> show paragon cluster details
Storage and Controller Node IPs: 10.1.2.3,10.1.2.4,10.1.2.5,10.1.2.6
```

Release Information

Command introduced in Juniper Paragon Automation Release 2.0.0.

show paragon cluster namespaces

IN THIS SECTION

- [Syntax | 1468](#)
- [Description | 1468](#)
- [Options | 1469](#)
- [Required Privilege Level | 1469](#)
- [Output Fields | 1469](#)
- [Sample Output | 1469](#)
- [Release Information | 1470](#)

Syntax

```
show paragon cluster namespaces
```

Description

Show namespace information of your Routing Director cluster.

Options

This command has no options.

Required Privilege Level

view

Output Fields

[Table 343 on page 1469](#) lists the output fields for the `show paragon cluster namespaces` command. Output fields are listed in the approximate order in which they appear.

Table 343: show paragon cluster namespaces Output Fields

Field Name	Field Description
NAME	Name of the namespace.
STATUS	Status of the namespace.
AGE	Time the namespace has been active.

Sample Output

`show paragon cluster namespaces`

```
user@node> show paragon cluster namespaces

NAME          STATUS  AGE
airflow       Active  11d
arango        Active  11d
calico-system Active  11d
cert-manager  Active  11d
common        Active  11d
cosign-system Active  11d
default       Active  11d
```

epic	Active	11d
foghorn	Active	11d
gnmi-term	Active	11d
healthbot	Active	11d
installer	Active	11d
kube-node-lease	Active	11d
kube-public	Active	11d
kube-system	Active	11d
metallb-system	Active	11d
netbox	Active	11d
northstar	Active	11d
oc-term	Active	11d
paa	Active	11d
papi	Active	11d
rook-ceph	Active	11d
streams	Active	11d
tigera-operator	Active	11d
traefik	Active	11d
traefik-paa	Active	11d
trust	Active	11d
victoria-metrics	Active	11d
webhooks	Active	11d

Release Information

Command introduced in Juniper Paragon Automation Release 2.0.0.

show paragon cluster nodes

IN THIS SECTION

- [Syntax | 1471](#)
- [Description | 1471](#)
- [Options | 1471](#)
- [Required Privilege Level | 1471](#)
- [Output Fields | 1471](#)

- [Sample Output | 1472](#)
- [Release Information | 1473](#)

Syntax

```
show paragon cluster nodes
```

```
show paragon cluster nodes head n sort (cpu|memory)
```

Description

Show node information of your Routing Director cluster.

Options

none	Display node information.
head <i>n</i>	(Optional) Display information about the first <i>n</i> nodes. Example: show paragon cluster nodes head 2 displays information about the first two nodes.
sort (cpu memory)	(Optional) Display information of the nodes sorted by cpu or memory utilization.

Required Privilege Level

view

Output Fields

[Table 344 on page 1472](#) lists the output fields for the show paragon cluster nodes command. Output fields are listed in the approximate order in which they appear.

Table 344: show paragon cluster nodes Output Fields

Field Name	Field Description
NAME	Name of the node.
STATUS	Status of the node.
ROLES	Role the node plays in the Kubernetes cluster.
AGE	Time the node has been active.
VERSION	Kubernetes version.

Sample Output

show paragon cluster nodes

```
user@node> show paragon cluster nodes
```

NAME	STATUS	ROLES	AGE	VERSION
node-0	Ready	control-plane,etcd,master	11d	v1.28.6+rke2r1
node-1	Ready	control-plane,etcd,master	11d	v1.28.6+rke2r1
node-2	Ready	control-plane,etcd,master	11d	v1.28.6+rke2r1
node-3	Ready	influxdb-worker,worker	11d	v1.28.6+rke2r1

show paragon cluster nodes head 2 sort cpu

```
user@node> show paragon cluster nodes head 2 sort cpu
```

NAME	CPU(cores)	CPU%	MEMORY(bytes)	MEMORY%
node-3	3017m	21%	15054Mi	52%
node-2	896m	6%	12813Mi	47%

Release Information

Command introduced in Juniper Paragon Automation Release 2.0.0.

show paragon cluster pods

IN THIS SECTION

- [Syntax | 1473](#)
- [Description | 1473](#)
- [Options | 1473](#)
- [Required Privilege Level | 1474](#)
- [Output Fields | 1474](#)
- [Sample Output | 1475](#)
- [Release Information | 1476](#)

Syntax

```
show paragon cluster pods
```

```
show paragon cluster pods namespace name head n sort (cpu|memory)
```

Description

Show information about pods in your Routing Director cluster.

Options

none	Display information about pods in the cluster.

head <i>n</i>	(Optional) Display information about the first <i>n</i> pods. Example: show paragon cluster pods head 3 displays information about the first three pods.
namespace <i>name</i>	(Optional) Display information of the pods in a namespace.
sort (cpu memory)	(Optional) Display information of the pods sorted by cpu or memory utilization.

Required Privilege Level

view

Output Fields

Table 345 on page 1474 lists the output fields for the show paragon cluster pods command. Output fields are listed in the approximate order in which they appear.

Table 345: show paragon cluster pods Output Fields

Field Name	Field Description
NAMESPACE	Name of the namespace.
NAME	Name of the pod.
READY	Number of pods that are ready.
STATUS	Status of the pod.
RESTARTS	Time since the pod was last restarted.
Age	Time the node has been active.

Sample Output

show paragon cluster pods

```
user@node> show paragon cluster pods
```

NAMESPACE	NAME		READY
STATUS	RESTARTS	AGE	
airflow	airflow-scheduler-9599489c6-75v6n		2/2
Running	1 (11d ago)	11d	
airflow	airflow-webserver-b4bfd554d-t7d89		1/1
Running	0	11d	
airflow	airflow-worker-7b6967654-pxxzc		1/1
Running	0	11d	
airflow	airflow-worker-7b6967654-rsl4h		1/1
Running	0	11d	
airflow	workflow-manager-7476b57f88-87g4z		1/1
Running	0	11d	
arango	arango-arango-operator-58dbbf6669-55bc7		1/1
Running	1 (11d ago)	11d	
arango	arango-arango-operator-58dbbf6669-wcn82		1/1
Running	1 (4d21h ago)	11d	
arango	foghorn-agnt-bsfcjebx-f450c7		1/1
Running	0	11d	
arango	foghorn-agnt-cunpycys-f450c7		1/1
Running	0	11d	

<output snipped>

show paragon cluster pods namespace airflow head 3 sort cpu

```
user@node> show paragon cluster pods namespace airflow head 3 sort cpu
```

NAME	CPU(cores)	MEMORY(bytes)
airflow-scheduler-9d86ffb65-gdn7h	1864m	737Mi
airflow-webserver-fdbdcf8-2plkj	3m	419Mi
airflow-worker-6c568dc6c8-n54lx	3m	1458Mi

Release Information

Command introduced in Juniper Paragon Automation Release 2.0.0.

show paragon cluster pods namespace healthbot sort

IN THIS SECTION

- [Syntax | 1476](#)
- [Description | 1476](#)
- [Options | 1476](#)
- [Required Privilege Level | 1477](#)
- [Output Fields | 1477](#)
- [Sample Output | 1477](#)
- [Release Information | 1478](#)

Syntax

```
show paragon cluster pods namespace healthbot sort (cpu|memory)
```

Description

Show pods of the healthbot namespace sorted by either CPU or memory utilization.

Options

sort (cpu memory)	Display information of the pods in the healthbot namespace sorted by cpu or memory utilization.

Required Privilege Level

view

Output Fields

[Table 346 on page 1477](#) lists the output fields for the `show paragon cluster pods namespace healthbot sort utilization type` command. Output fields are listed in the approximate order in which they appear.

Table 346: show paragon cluster pods namespace healthbot sort *utilization type* Output Fields

Field Name	Field Description
NAME	Name of the pod in the healthbot namespace.
CPU(crores)	CPU utilization of the pod.
MEMORY(bytes)	Memory utilization of the pod.

Sample Output

show paragon cluster pods namespace healthbot sort cpu

```

user@node> show paragon cluster pods namespace healthbot sort cpu
NAME                                CPU(crores)  MEMORY(bytes)
alerta-74dbdb998d-v7zs9             55m          771Mi
pyez-scheduler-7745f56b9b-mqbbw     26m          118Mi
pyez-scheduler-7745f56b9b-smrw6     25m          136Mi
alerta-74dbdb998d-2hqrm             22m          757Mi
pyez-scheduler-7745f56b9b-6v9vn     19m          121Mi
config-server-7d99d65874-8bddq      16m          338Mi
analytical-engine-799bb69799-rhsr4  14m          24Mi
analytical-engine-799bb69799-ltb7h  12m          19Mi
analytical-engine-799bb69799-xwn5t  12m          20Mi
itsdb-677b4fcb89-mdqm8              11m          32Mi
itsdb-677b4fcb89-9scds              10m          35Mi
tand-59ffb8cf98-ndn7f               10m          21Mi
tand-59ffb8cf98-6jpsl                7m          13Mi

```

tand-59ffb8cf98-np7wb	7m	16Mi
jtimon-5645f59bc6-k6xlx	5m	46Mi
jtimon-5645f59bc6-5fcqn	4m	39Mi
tsdb-shim-648d49c7f7-64z99	4m	20Mi
tsdb-shim-648d49c7f7-9hgcf	4m	25Mi
tsdb-shim-648d49c7f7-grkd6	4m	18Mi
influxdb-influxdb1-84b75fc5b6-cgtbk	1m	28Mi
api-server-7cb454d9dd-sg4wk	1m	134Mi
api-server-7cb454d9dd-n7558	1m	129Mi

show paragon cluster pods namespace healthbot sort memory

```
user@node> show paragon cluster pods namespace healthbot sort memory
```

NAME	CPU(cores)	MEMORY(bytes)
alerta-74dbdb998d-v7zs9	30m	771Mi
alerta-74dbdb998d-2hqrm	35m	757Mi
config-server-7d99d65874-8bddq	12m	338Mi
pyez-scheduler-7745f56b9b-smrw6	23m	136Mi
api-server-7cb454d9dd-sg4wk	1m	134Mi
api-server-7cb454d9dd-n7558	1m	129Mi
pyez-scheduler-7745f56b9b-6v9vn	10m	121Mi
pyez-scheduler-7745f56b9b-mqbbw	15m	118Mi
jtimon-5645f59bc6-k6xlx	5m	45Mi
jtimon-5645f59bc6-5fcqn	4m	39Mi
itsdb-677b4fcb89-9scds	10m	35Mi
itsdb-677b4fcb89-mdqm8	11m	33Mi
influxdb-influxdb1-84b75fc5b6-cgtbk	1m	28Mi
tsdb-shim-648d49c7f7-9hgcf	4m	26Mi
analytical-engine-799bb69799-rhsr4	14m	24Mi
tand-59ffb8cf98-ndn7f	9m	21Mi
tsdb-shim-648d49c7f7-64z99	5m	20Mi
analytical-engine-799bb69799-xwn5t	12m	20Mi
analytical-engine-799bb69799-ltb7h	13m	19Mi
tsdb-shim-648d49c7f7-grkd6	5m	18Mi
tand-59ffb8cf98-np7wb	7m	16Mi
tand-59ffb8cf98-6jpsl	8m	13Mi

Release Information

Command introduced in Juniper Paragon Automation Release 2.0.0.

show paragon images version

IN THIS SECTION

- [Syntax | 1479](#)
- [Description | 1479](#)
- [Options | 1479](#)
- [Required Privilege Level | 1479](#)
- [Output Fields | 1480](#)
- [Sample Output | 1480](#)
- [Release Information | 1481](#)

Syntax

```
show paragon images version

show paragon images version namespace namespace-name
```

Description

Show the version of pods of your Routing Director cluster.

Options

none	Display version of pods in the cluster.
"namespace namespace-name" on page 1481	Display image version of a namespace.

Required Privilege Level

view

Output Fields

Table 347 on page 1480 lists the output fields for the `show paragon images version` command. Output fields are listed in the approximate order in which they appear.

Table 347: `show paragon images version` Output Fields

Field Name	Field Description
Namespace	Name of the namespace.
Pod	Name of a pod within a namespace.
Image	Version of a pod within a namespace.

Sample Output

`show paragon images version`

```
user@node> show paragon images version

[Summary]: 24 namespaces, 222 pods, 357 containers and 125 different images
+-----+-----+
+-----+-----+
+-----+-----+
| Namespace | Pod |
|          |     |
Image |
+-----+-----+
+-----+-----+
+-----+-----+
|          |     |
|          |     |
|          |     |
+-----+-----+
+-----+-----+
+-----+-----+
| airflow | airflow-scheduler-7765648756-xqp26 | paragon-
registry.local/northstar-docker-local/
```

```
airflow2:0.1.1
|
+           +
+-----+
-----+
|           |                               | paragon-
registry.local/airflow2/workflow-
manager:1.119.0
|
+           +-----+
+-----+
-----+
|           | airflow-webserver-f7647b7c9-rds69           | paragon-
registry.local/northstar-docker-local/airflow2:0.1.1

<Output Snipped>
```

Release Information

Command introduced in Juniper Paragon Automation Release 2.0.0.

show paragon images version namespace

IN THIS SECTION

- [Syntax | 1482](#)
- [Description | 1482](#)
- [Options | 1482](#)
- [Required Privilege Level | 1482](#)
- [Output Fields | 1482](#)
- [Sample Output | 1483](#)
- [Release Information | 1484](#)

Syntax

```
show paragon images version namespace namespace-name
```

Description

Show the image version of specific namespaces of your Routing Director setup.

Options

namespace healthbot	Display the image version of the healthbot namespace.
namespace epic	Display the image version of the epic namespace.
namespace foghorn	Display the image version of the foghorn namespace.
namespace northstar	Display the image version of the northstar namespace.
namespace papi	Display the image version of the papi namespace.

Required Privilege Level

view

Output Fields

"[show paragon images version namespace](#)" on [page 1481](#) lists the output fields for the show paragon cluster nodes command. Output fields are listed in the approximate order in which they appear.

Table 348: show paragon cluster nodes Output Fields

Field Name	Field Description
NAMESPACE	Name of the namespace.

Table 348: show paragon cluster nodes Output Fields *(Continued)*

Field Name	Field Description
POD	Name of a pod within the namespace.
IMAGE	Version of a pod within the namespace.

Sample Output

show paragon images version namespace

```

user@node> show paragon images version namespace healthbot
[Summary]: 1 namespaces, 24 pods, 39 containers and 12 different images
+-----+-----+
+-----+-----+
+-----+-----+
| Namespace | Pod
|
Image |
+-----+-----+
+-----+-----+
+-----+-----+
| healthbot | alerta-74dbdb998d-2hqrm | paragon-registry.local/insights/alerta@sha256:615128a72487fa58ee21ca48225bce10cccb37547f908c85e8e792ea4ededc3f |
+ +
+-----+-----+
+-----+-----+
| | | paragon-registry.local/insights/init-container@sha256:2b5110f19c794ea6dfc3995f239bb18eb3d2d66b23a148687769c94fa23a509e |
+ +-----+-----+
+-----+-----+
+-----+-----+
| | alerta-74dbdb998d-v7zs9 | paragon-registry.local/insights/alerta@sha256:615128a72487fa58ee21ca48225bce10cccb37547f908c85e8e792ea4ededc3f |
+ +
+-----+-----+
+-----+-----+
| | | paragon-registry.local/insights/init-

```

```

container@sha256:2b5110f19c794ea6dfc3995f239bb18eb3d2d66b23a148687769c94fa23a509e |
+      +-----+
+-----+
+-----+
|      | analytical-engine-799bb69799-ltb7h | paragon-registry.local/insights/analytical-
engine@sha256:d1858db9ba06da7b5dba4b84df0ce06e0a4831f07ec306131988f7e56350d30c |
+      +-----+
+
+
+
|      | analytical-engine-799bb69799-rhsr4
|
|
+      +-----+
+
+
+
|      | analytical-engine-799bb69799-xwn5t
|
|
+      +-----+
+-----+
+-----+
|      | api-server-7cb454d9dd-n7558      | paragon-registry.local/insights/api-
server@sha256:62a088a728e3b7a1ee48667d388e7ed4c4e2a709d0b1c932b1ff7a95a4373de4

<Output Snipped>

```

Release Information

Command introduced in Juniper Paragon Automation Release 2.0.0.

show paragon pvc details

IN THIS SECTION

- [Syntax | 1485](#)
- [Description | 1485](#)
- [Options | 1485](#)

- [Required Privilege Level | 1485](#)
- [Output Fields | 1485](#)
- [Sample Output | 1487](#)
- [Release Information | 1489](#)

Syntax

```
show paragon pvc details
```

Description

Show persistent volume (PV) and persistent volume claim (PVC) information.

Options

This command has no options.

Required Privilege Level

view

Output Fields

[Table 349 on page 1485](#) lists the output fields for the `show paragon pvc details` command. Output fields are listed in the approximate order in which they appear.

Table 349: show paragon pvc details Output Fields

Field Name	Field Description
PV details	
NAME	Name of the persistent volume.

Table 349: show paragon pvc details Output Fields (Continued)

Field Name	Field Description
CAPACITY	Storage capacity of the persistent volume.
ACCESS MODES	Access modes supported by the persistent volume.
RECLAIM POLICY	Reclaim policy for the persistent volume.
STATUS	Current status of the persistent volume.
CLAIM	Name of the persistent volume claim bound to the persistent volume.
STORAGECLASS	Storage class, if any, used by the persistent volume.
REASON	Reason for the current state of the persistent volume.
AGE	Total time the persistent volume has been in the current state.
PVC details	
NAMESPACE	Namespace that the persistent volume claim is a part of.
NAME	Name of the persistent volume claim.
STATUS	Current state of the persistent volume claim.
VOLUME	Name of the persistent volume bound to the persistent volume claim.

Table 349: show paragon pvc details Output Fields *(Continued)*

Field Name	Field Description
CAPACITY	Storage capacity requested by the persistent volume claim.
ACCESS MODE	Access modes requested by the persistent volume claim.
STORAGECLASS	Storage class used by the persistent volume claim.
AGE	Total time the persistent volume claim has been in the current state.

Sample Output

show paragon pvc details

```
user@node> show paragon pvc details
```

PV details						
NAME			CAPACITY	ACCESS MODES	RECLAIM	POLICY
STATUS	CLAIM					
STORAGECLASS	REASON	AGE				
local-pv-1126409			97Gi	RWO	Delete	
Bound	arango/foghorn-agent-bsfcjebx					local-
storage		11d				
local-pv-175fd7c6			97Gi	RWO	Delete	
Available						local-
storage		11d				
local-pv-229f0f1b			97Gi	RWO	Delete	
Available						local-
storage		11d				
local-pv-2c329a72			97Gi	RWO	Delete	
Available						local-
storage		11d				
local-pv-409bac0b			97Gi	RWO	Delete	

Bound	common/data-zookeeper-2				local-
storage	11d				
local-pv-417ecfef		97Gi	RWO	Delete	
Bound	common/pgdata-atom-db-2				local-
storage	11d				
local-pv-5452ea4f		97Gi	RWO	Delete	
Available					local-
storage	11d				
local-pv-5f1f9ecf		97Gi	RWO	Delete	
Available					local-
storage	11d				
local-pv-6e3f1422		97Gi	RWO	Delete	
Bound	arango/foghorn-dbserver-wmbr3ond				local-
storage	11d				
local-pv-7181e6e5		97Gi	RWO	Delete	
Available					local-
storage	11d				
local-pv-72423f89		97Gi	RWO	Delete	
Available					local-
storage	11d				

PVC details

NAMESPACE	NAME			STATUS	
VOLUME		CAPACITY	ACCESS MODES	STORAGECLASS	AGE
airflow	airflow			Bound	
pvc-4ed0d0b7-96cc-4d74-8188-af9950051606		10Gi	RWX	rook-cephfs	11d
arango	foghorn-agent-bsfcjebx			Bound	local-
pv-1126409		97Gi	RWO	local-storage	11d
arango	foghorn-agent-cunpycys			Bound	local-pv-
eeff4bc8		97Gi	RWO	local-storage	11d
arango	foghorn-agent-vfhd4ydv			Bound	local-pv-
c932346		97Gi	RWO	local-storage	11d
arango	foghorn-dbserver-6h8se2vq			Bound	local-
pv-889c24c8		97Gi	RWO	local-storage	11d
arango	foghorn-dbserver-o8tstx85			Bound	local-pv-
a35f649b		97Gi	RWO	local-storage	11d
arango	foghorn-dbserver-wmbr3ond			Bound	local-
pv-6e3f1422		97Gi	RWO	local-storage	11d
common	data-kafka-0			Bound	local-pv-
df00ed4		97Gi	RWO	local-storage	11d
common	data-kafka-1			Bound	local-pv-
f2423c65		97Gi	RWO	local-storage	11d

Release Information

Command introduced in Juniper Paragon Automation Release 2.0.0.

show paragon version

IN THIS SECTION

- [Syntax | 1489](#)
- [Description | 1489](#)
- [Options | 1489](#)
- [Required Privilege Level | 1489](#)
- [Output Fields | 1490](#)
- [Sample Output | 1490](#)
- [Release Information | 1490](#)

Syntax

```
show paragon version
```

Description

Show version information of your Routing Director cluster.

Options

This command has no options.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided with the build and Kubernetes version information.

Sample Output

show paragon version

```
user@node> show paragon version

ova: 20240403_2103
build: sample-release-2.0.0.6802.g5f4b66517d
Client Version: v1.28.6+rke2r1
Kustomize Version: v5.0.4-0.20230601165947-6ce0bf390ce3
Server Version: v1.28.6+rke2r1
```

Release Information

Command introduced in Juniper Paragon Automation Release 2.0.0.

Configuration Mode Commands

IN THIS CHAPTER

- [delete paragon cluster | 1491](#)
- [load set | 1497](#)
- [set paragon cluster applications | 1498](#)
- [set paragon cluster common-services | 1504](#)
- [set paragon cluster insights victoria-metrics | 1507](#)
- [set paragon cluster install | 1509](#)
- [set paragon cluster kubernetes | 1511](#)
- [set paragon cluster mail-server | 1512](#)
- [set paragon cluster nodes | 1515](#)
- [set paragon cluster ntp | 1516](#)
- [set paragon cluster papi | 1518](#)
- [set paragon monitoring | 1519](#)
- [set system login | 1522](#)
- [show paragon cluster \(c\) | 1524](#)
- [show paragon cluster applications \(c\) | 1529](#)
- [show paragon cluster common-services \(c\) | 1532](#)

delete paragon cluster

IN THIS SECTION

- [Syntax | 1492](#)
- [Hierarchy Level | 1492](#)
- [Description | 1493](#)

- [Options | 1493](#)
- [Required Privilege Level | 1497](#)
- [Release Information | 1497](#)

Syntax

```
delete paragon cluster applications (none|applications_options)
```

```
delete paragon cluster common-services (none|common_services_options)
```

```
delete paragon cluster install (none|offline-install|prefer-ipv6)
```

```
delete configuration paragon cluster kubernetes (none|address-family [none|cluster-ipv6-enabled])
```

```
delete paragon cluster mail-server (none|mail_server_options)
```

```
delete paragon cluster nodes (none|nodes_options)
```

```
delete paragon cluster ntp (none|ntp-servers)
```

```
delete paragon cluster papi (none|papi-local-user-management)
```

Hierarchy Level

[edit]

Description

Delete the specified parameters under the paragon hierarchy.

Options

applications (none <i>applications_options</i>)	Delete the configured applications settings. <i>applications_options</i> options include: • active-assurance • active-assurance test-agent-gateway-vip • active-assurance test-agent-gateway-vip-ipv6 • active-assurance test-agent-gateway-hostname • aiops • aiops enable-device-health • aiops install-aiml • gnmi-term • gnmi-term gnmi-term-host • gnmi-term gnmi-term-host-ipv6 • oc-term • oc-term oc-term-host • oc-term oc-term-host-ipv6 • oc-term oc-term-previous-management-ip-family • pathfinder pce-server pce-server-vip • pathfinder pce-server pce-server-global-default-tls-mode • pathfinder pce-server pce-server-user-certificate pce-server-use-custom-certificate • pathfinder pce-server pce-server-user-certificate pce-server-custom-certificate-filename • pathfinder pce-server pce-server-user-certificate pce-server-custom-certificate-key-filename • routingbot

	• routingbot install-routingbot • routingbot routingbot-crpd-vip • web-ui • web-ui web-admin-password • web-ui web-admin-user For more information, see "set paragon cluster applications" on page 1498.
common-services (none <i>common_services_options</i>)	Delete the common services settings. <i>common_services_options</i> options include: • ingress ingress-hostname • ingress ingress-vip • ingress ingress-vip-ipv6 • ingress user-certificate • ingress user-certificate use-user-certificate • ingress user-certificate user-certificate-filename • ingress user-certificate-key-filename • metallb metallb-bgp-peer peer-ip • metallb metallb-bgp-peer peer-asn • metallb metallb-bgp-peer local-asn • metallb metallb-bgp-peer local-nodes For more information, see "set paragon cluster common-services" on page 1504.
install (none offline-install [none true false] prefer-ipv6 [none true false])	Delete the cluster installation type setting. For more information, see "set paragon cluster install" on page 1509.

kubernetes (none address-family [none cluster-ipv6-enabled {none true false}])	Delete the usage of IPv6 addresses for the cluster. For more information, see "set paragon cluster kubernetes" on page 1511.
mail-server (none <i>mail_server_options</i>)	Delete SMTP configuration settings. <i>mail_server_options</i> options include: • smtp-allowed-sender-domains • smtp-relayhost • smtp-relayhost-password • smtp-relayhost-username • smtp-sender-email • smtp-sender-name For more information, see "set paragon cluster mail-server" on page 1512.
nodes (none <i>nodes_options</i>)	Delete the configured node parameters. <i>nodes_options</i> options include: • kubernetes <i>index</i> • kubernetes <i>index</i> address • kubernetes <i>index</i> hostname For more information, see "set paragon cluster nodes" on page 1515.
ntp (none ntp-servers)	Delete the configured NTP servers. For more information, see "set paragon cluster ntp" on page 1516.
papi (none papi-local-user-management)	Delete the local user management on the nodes setting. For more information, see "set paragon cluster papi" on page 1518.

Required Privilege Level

configure

Release Information

Statement introduced in Juniper Paragon Automation Release 2.0.0.

pathfinder option introduced in Juniper Paragon Automation Release 2.1.0.

kubernetes, prefer-ipv6, ingress ingress-vip-ipv6, and test-agent-gateway-vip-ipv6 options introduced in Juniper Paragon Automation Release 2.2.0.

aiops, gnmi-term, oc-term, routingbot, and metallb options introduced in Juniper Paragon Automation Release 2.4.0.

pce-server-global-default-tls-mode and pce-server-user-certificate options introduced in Juniper Routing Director Release 2.5.0.

RELATED DOCUMENTATION

| [show configuration paragon cluster](#) | [1439](#)

load set

IN THIS SECTION

- [Syntax](#) | [1498](#)
- [Hierarchy Level](#) | [1498](#)
- [Description](#) | [1498](#)
- [Options](#) | [1498](#)
- [Required Privilege Level](#) | [1498](#)
- [Release Information](#) | [1498](#)

Syntax

```
load set filename
```

Hierarchy Level

[edit]

Description

Load a list of set commands configuration saved in an ASCII file. You must commit the configuration after loading to activate it.

Options

<i>filename</i>	Enter the filename which contains the information to be configured. The file can be remote or local.

Required Privilege Level

configure

Release Information

Command introduced in Juniper Paragon Automation Release 2.0.0.

set paragon cluster applications

IN THIS SECTION

- [Syntax | 1500](#)
- [Hierarchy Level | 1501](#)
- [Description | 1501](#)

- Options | 1502
- Required Privilege Level | 1504
- Release Information | 1504

Syntax

```
set paragon cluster applications active-assurance test-agent-gateway-vip v4address
```

```
set paragon cluster applications active-assurance test-agent-gateway-vip-ipv6 v6address
```

```
set paragon cluster applications active-assurance test-agent-gateway-hostname hostname
```

```
set paragon cluster applications aiops enable-device-health (true|false)
```

```
set paragon cluster applications aiops install-aiml (true|false)
```

```
set paragon cluster applications gnmi-term gnmi-term-host v4address-or-hostname
```

```
set paragon cluster applications gnmi-term gnmi-term-host-ipv6 v6address-or-hostname
```

```
set paragon cluster applications oc-term oc-term-host v4address-or-hostname
```

```
set paragon cluster applications oc-term oc-term-host-ipv6 v6address-or-hostname
```

```
set paragon cluster applications pathfinder pce-server pce-server-vip pce-server-vip
```

```
set paragon cluster applications pathfinder pce-server pce-server-global-default-tls-mode (auto-detect|strict-enable|strict-disable)
```

```
set paragon cluster applications pathfinder pce-server pce-server-user-certificate pce-server-use-custom-certificate (true|false)
```

```
set paragon cluster applications pathfinder pce-server pce-server-user-certificate
```

```
pce-server-custom-certificate-filename cert-filename
```

```
set paragon cluster applications pathfinder pce-server pce-server-usercertificate  
pce-server-custom-certificate-key-filename privatekey-filename
```

```
set paragon cluster applications routingbot install-routingbot (true|false)
```

```
set paragon cluster applications routingbot routingbot-crpd-vip v4-address
```

```
set paragon cluster applications web-ui web-admin-password password
```

```
set paragon cluster applications web-ui web-admin-user user_e-mail_id
```

Hierarchy Level

[edit]

Description

Configure the following application parameters:

- Paragon Active Assurance Test Agent gateway virtual IP (VIP) address and hostname
- Default Web administrator login credentials. The Web administrator is the first user to log in to the Routing Director GUI and has super user privileges.
- PCE server VIP address used to set up Path Computational Element Protocol (PCEP) sessions between Routing Director and the devices.
- Routing observability use case and VIP address.
- Enable AI/ML (artificial intelligence [AI] and machine learning [ML]) techniques to automatically monitor Key Performance Indicators (KPIs) related to a device's health.
- Separate IP addresses for gNMI and NETCONF. If not configured, gNMI uses the generic ingress VIP address which is shared between Web GUI, gNMI, and NETCONF.

Options

active-assurance test-agent-gateway-vip <i>v4address</i>	Configure the IPv4 VIP address for the Paragon Active Assurance Test Agent gateway.
active-assurance test-agent-gateway-vip-ipv6 <i>v6address</i>	Configure the IPv6 VIP address for the Paragon Active Assurance Test Agent gateway.
active-assurance test-agent-gateway-hostname <i>hostname</i>	Configure the hostname for the Paragon Active Assurance Test Agent gateway.
aiops enable-device-health (true false)	Configure monitoring of device-health.
aiops install-aiml (true false)	Enable or disable AI/ML and related features. This feature is disabled by default.
gnmi-term gnmi-term-host <i>v4address-or-hostname</i>	Configure a separate VIP address or hostname for gNMI. If not configured, gNMI uses the generic ingress VIP address which is shared between Web GUI, gNMI, and NETCONF.
gnmi-term gnmi-term-host-ipv6 <i>v6address-or-hostname</i>	Configure a separate virtual IPv6 address or hostname for gNMI.
oc-term oc-term-host <i>v4address-or-hostname</i>	Configure a separate VIP address or hostname for NETCONF. If not configured, NETCONF uses the generic ingress VIP address which is shared between Web GUI, gNMI, and NETCONF.
oc-term oc-term-host-ipv6 <i>v6address-or-hostname</i>	Configure a separate virtual IPv6 address or hostname for NETCONF.
pathfinder pce-server pce-server-vip <i>pce-server-vip</i>	Configure the PCE Server VIP address.

pathfinder pce-server pce-server-global-default-tls-mode (auto-detect strict-enable strict-disable)	Enable security between the PCE server and Path Computation Clients (PCC). Options include: • auto-detect—Accepts PCEP connection regardless of whether security is enabled. • strict-enable—Enforces security on all PCEP sessions. If a device does not have security enabled, its connection is rejected. • strict-disable—Accepts only non-secured PCEP connections.
pathfinder pce-server pce-server-user-certificate pce-server-use-custom-certificate (true false)	Enable usage of custom certificates to validate secure PCEP connections between the device and Routing Director.
pathfinder pce-server pce-server-user-certificate pce-server-custom-certificate-filename <i>cert-filename</i>	Configure the custom certificate file used to validate secure PCEP connections between the device and Routing Director.
pathfinder pce-server pce-server-user-certificate pce-server-custom-certificate-key-filename <i>privatekey-filename</i>	Configure the custom certificate key used to validate secure PCEP connections between the device and Routing Director.
routingbot install-routingbot (true false)	Enable or disable the routing observability feature to monitor routing health in real-time. This feature is enabled by default.
routingbot routingbot-crpd-vip <i>v4-address</i>	VIP to connect to cRPD for routing observability.
web-ui web-admin-password <i>password</i>	Configure the password for the Web administrator. Ensure that the password does not start with 'ENC:'.
web-ui web-admin-user <i>user_e-mail_id</i>	Configure the login e-mail ID for the Web administrator.

Required Privilege Level

configure

Release Information

Statement introduced in Juniper Paragon Automation Release 2.0.0.

pathfinder option introduced in Juniper Paragon Automation Release 2.1.0.

test-agent-gateway-vip-ipv6 option introduced in Juniper Paragon Automation Release 2.2.0.

aiops, gnmi-term, oc-term, and routingbot options introduced in Juniper Paragon Automation Release 2.4.0.

pce-server-global-default-tls-mode and pce-server-user-certificate options introduced in Juniper Routing Director Release 2.5.0.

RELATED DOCUMENTATION

[show configuration paragon cluster | 1439](#)

[delete paragon cluster | 1491](#)

set paragon cluster common-services

IN THIS SECTION

- [Syntax | 1505](#)
- [Hierarchy Level | 1505](#)
- [Description | 1505](#)
- [Options | 1505](#)
- [Required Privilege Level | 1507](#)
- [Release Information | 1507](#)

Syntax

```
set paragon cluster common-services ingress ingress-vip v4address
```

```
set paragon cluster common-services ingress ingress-vip-ipv6 v6address
```

```
set paragon cluster common-services ingress ingress-hostname hostname
```

```
set paragon cluster common-services ingress user-certificate use-user-certificate (true|false)  
user-certificate-filename cert_file user-certificate-key-filename cert_key_file
```

```
set paragon cluster common-services metallb metallb-bgp-peer peer-ip ToR-v4address peer-asn ToR-  
asn local-asn node-asn local-nodes (node-IP [node1-IP node2-IP ...])
```

```
set paragon cluster common-services metallb metallb-bgp-peer-ipv6 peer-ip ToR-v6address peer-asn  
ToR-asn local-asn node-asn local-nodes (node-IPv6 [node1-IPv6 node2-IPv6 ...])
```

Hierarchy Level

[edit]

Description

Configure the common ingress virtual IP (VIP) address and hostname. The ingress VIP is used to access the Paragon Automation Web GUI. Also, enable or disable custom user certificates. If enabled, the certificate file and key file must be copied to the **/root/epic/config** directory.

If your cluster nodes are in different subnets, configure BGP peering between the top-of-rack (ToR) router and the cluster nodes using the metalLB agent running in each cluster node.

Options

ingress-hostname <i>hostname</i>	Configure the hostname for the ingress VIP address.
ingress-vip <i>v4address</i>	Configure the common ingress IPv4 VIP address.
ingress-vip-ipv6 <i>v6address</i>	Configure the common ingress IPv6 VIP address.
user-certificate use-user-certificate (true false)	Enable or disable custom user certificates.
user-certificate user-certificate-filename <i>cert_file</i>	Configure the custom user certificate filename.
user-certificate user-certificate-key-filename <i>cert_key_file</i>	Configure the custom user certificate key filename.
metallb metallb-bgp-peer peer-ip <i>ToR-v4address</i> peer-asn <i>ToR-asn</i> local-asn <i>node-asn</i> local-nodes (<i>node-IP</i> [<i>node1-IP node2-IP ...</i>])	Configure BGP peering between the ToR and the cluster nodes using the metalLB agent running in each cluster node. • peer-ip is the IP address of the ToR router • peer-asn is the AS number of the ToR router • local-asn is AS number of the cluster nodes • local-nodes is cluster node IP address. You can enter one or more addresses.
metallb metallb-bgp-peer-ipv6 peer-ip <i>ToR-v6address</i> peer-asn <i>ToR-asn</i> local-asn <i>node-asn</i> local-nodes (<i>node-IPv6</i> [<i>node1-IPv6 node2-IPv6 ...</i>])	Configure BGP peering between the ToR and the cluster nodes using the metalLB agent running in each cluster node. • peer-ip is the IPv6 address of the ToR router • peer-asn is the AS number of the ToR router • local-asn is AS number of the cluster nodes • local-nodes is cluster node IPv6 address. You can enter one or more addresses.

Required Privilege Level

configure

Release Information

Statement introduced in Juniper Paragon Automation Release 2.0.0.

ingress-vip-ipv6 option added in Juniper Paragon Automation Release 2.2.0.

metallb option added in Juniper Paragon Automation Release 2.4.0.

RELATED DOCUMENTATION

[show configuration paragon cluster | 1439](#)

[delete paragon cluster | 1491](#)

set paragon cluster insights victoria-metrics

IN THIS SECTION

- [Syntax | 1508](#)
- [Hierarchy Level | 1508](#)
- [Description | 1508](#)
- [Options | 1508](#)
- [Required Privilege Level | 1508](#)
- [Release Information | 1509](#)

Syntax

```
set paragon cluster insights victoria-metrics vm-cluster-replication-factor replication_factor
```

```
set paragon cluster insights victoria-metrics vm-cluster-retention-period retention_period
```

```
set paragon cluster insights victoria-metrics vm-cluster-storage-instances number_of_pods
```

Hierarchy Level

[edit]

Description

Update the Victoria Metrics cluster configuration parameters.

If you are updating the parameters post installation of the Routing Director cluster, use the request `paragon deploy cluster input "-v -t victoriametrics-cluster"` command to deploy the updates. For more information, see *Post Installation Tasks*.

Options

vm-cluster-replication-factor <i>replication_factor</i>	Set the Victoria Metrics cluster replication factor. Default: 1.
vm-cluster-retention-period <i>retention_period</i>	Set the Victoria Metrics retention period. The period should be in the [1-9][0-9]*[dmy] pattern. Default: 7days.
vm-cluster-storage-instances <i>number_of_pods</i>	Set the number of storage pods in the Victoria Metrics cluster. Default: 3.

Required Privilege Level

configure

Release Information

Statement introduced in Juniper Paragon Automation Release 2.1.0.

RELATED DOCUMENTATION

[show configuration paragon cluster | 1439](#)

[request paragon deploy cluster | 1403](#)

set paragon cluster install

IN THIS SECTION

- [Syntax | 1509](#)
- [Hierarchy Level | 1510](#)
- [Description | 1510](#)
- [Options | 1510](#)
- [Required Privilege Level | 1510](#)
- [Release Information | 1510](#)

Syntax

```
set paragon cluster install offline-install (true|false)
```

```
set paragon cluster install prefer-ipv6 (true|false)
```

```
set paragon cluster install enable-l3-vip (true|false)
```

Hierarchy Level

[edit]

Description

Enable or disable air-gap installation of the Paragon Automation cluster or usage of IPv6 addresses over IPv4 for the cluster configuration. Enable configuration of cluster nodes and VIP addresses from multiple subnets.

Options

offline-install (true false)	Set to true to enable air-gap installation or false to disable it.
prefer-ipv6 (true false)	Use IPv6 addresses over IPv4 addresses for cluster configuration.
enable-l3-vip (true false)	Configure IP addresses of cluster nodes and VIP addresses from multiple subnets.

Required Privilege Level

configure

Release Information

Statement introduced in Juniper Paragon Automation Release 2.0.0.

prefer-ipv6 option introduced in Juniper Paragon Automation Release 2.2.0.

enable-l3-vip option introduced in Juniper Paragon Automation Release 2.4.0.

RELATED DOCUMENTATION

show configuration paragon cluster 1439
delete paragon cluster 1491

set paragon cluster kubernetes

IN THIS SECTION

- [Syntax | 1511](#)
- [Hierarchy Level | 1511](#)
- [Description | 1511](#)
- [Options | 1511](#)
- [Required Privilege Level | 1511](#)
- [Release Information | 1512](#)

Syntax

```
set paragon cluster kubernetes address-family cluster-ipv6-enabled (true|false)
```

Hierarchy Level

[edit]

Description

Enable usage of IPv6 addresses for the Paragon Automation cluster.

Options

address-family cluster-ipv6-enabled (true false)	Enable or disable usage of IPv6 addresses for the cluster.

Required Privilege Level

configure

Release Information

Statement introduced in Juniper Paragon Automation Release 2.2.0.

RELATED DOCUMENTATION

[show configuration paragon cluster | 1439](#)

[delete paragon cluster | 1491](#)

set paragon cluster mail-server

IN THIS SECTION

- [Syntax | 1513](#)
- [Hierarchy Level | 1513](#)
- [Description | 1513](#)
- [Options | 1513](#)
- [Required Privilege Level | 1514](#)
- [Release Information | 1514](#)

Syntax

```
set paragon cluster mail-server smtp-allowed-sender-domains sender_domains
```

```
set paragon cluster mail-server smtp-relayhost hostname
```

```
set paragon cluster mail-server smtp-relayhost-password relayhost_password
```

```
set paragon cluster mail-server smtp-relayhost-username relayhost_username
```

```
set paragon cluster mail-server smtp-sender-email sender_email_address
```

```
set paragon cluster mail-server smtp-sender-name sender_name
```

```
set paragon cluster mail-server smtp-enabled (false|true)
```

Hierarchy Level

[edit]

Description

Configure SMTP to notify users when their accounts are created, activated, or modified.

If you are updating the SMTP settings post installation of the Routing Director cluster, use the request `paragon deploy cluster input "-t addon-apps,papi -e target_components=papi-kafka-emitter,mailservice"` command to deploy the updates. For more information, see [Configure SMTP Settings in Paragon Shell](#).

Options

smtp-allowed-sender-domains <i>sender_domains</i>	Set the e-mail domains from which Routing Director sends e-mails to users.
smtp-relayhost <i>hostname</i>	Set the name of the SMTP server that relays messages
smtp-relayhost-password <i>relayhost_password</i>	(Optional) Set the password for the SMTP (relay) server.
smtp-relayhost-username <i>relayhost_username</i>	(Optional) Set the user name to access the SMTP (relay) server.
smtp-sender-email <i>sender_email_address</i>	Set the e-mail address that appears as the sender's e-mail address to the e-mail recipient.
smtp-sender-name <i>sender_name</i>	Set the name that appears as the sender's name in the e-mails sent to users from Routing Director.
smtp-enabled (false true)	(Mandatory) Enable or disable SMTP.

Required Privilege Level

configure

Release Information

Statement introduced in Juniper Paragon Automation Release 2.0.0.

smtp-enabled option introduced in Juniper Paragon Automation Release 2.1.0.

RELATED DOCUMENTATION

[set paragon cluster papi | 1518](#)

[show configuration paragon cluster | 1439](#)

[delete paragon cluster | 1491](#)

[Configure SMTP Settings in Paragon Shell | 101](#)

Post Installation Tasks

set paragon cluster nodes

IN THIS SECTION

- [Syntax | 1515](#)
- [Hierarchy Level | 1515](#)
- [Description | 1515](#)
- [Options | 1515](#)
- [Required Privilege Level | 1516](#)
- [Release Information | 1516](#)

Syntax

```
set paragon cluster nodes kubernetes index address node-IP-address
```

```
set paragon cluster nodes kubernetes index hostname node-hostname
```

Hierarchy Level

[edit]

Description

Set the IP address or hostname of the Routing Director cluster node and configure the index number (1 through 4) of the nodes. The index number differentiates the order of the nodes in the cluster. 1 indicates the installer primary and worker node and 4 indicates the worker-only node.

Options

kubernetes <i>index</i>	(Mandatory) Set the Kubernetes node index number of the cluster node.

address <i>node_IP_address</i>	Set the interface IP address of the cluster node.
hostname <i>node_hostname</i>	Set the hostname of the cluster node.

Required Privilege Level

configure

Release Information

Statement introduced in Juniper Paragon Automation Release 2.0.0.

RELATED DOCUMENTATION

show configuration paragon cluster 1439
delete paragon cluster 1491

set paragon cluster ntp

IN THIS SECTION

- [Syntax | 1517](#)
- [Hierarchy Level | 1517](#)
- [Description | 1517](#)
- [Options | 1517](#)
- [Required Privilege Level | 1517](#)
- [Release Information | 1517](#)

Syntax

```
set paragon cluster ntp ntp-servers (server_name [set_of_ntp_servers])
```

Hierarchy Level

[edit]

Description

Configure NTP on the cluster nodes.

Options

ntp-servers <i>server_name</i>	Set the NTP server to sync to.
ntp-servers [<i>set_of_ntp_servers</i>]	Set multiple servers to sync to.

Required Privilege Level

configure

Release Information

Statement introduced in Juniper Paragon Automation Release 2.0.0.

RELATED DOCUMENTATION

show configuration paragon cluster 1439
delete paragon cluster 1491

set paragon cluster papi

IN THIS SECTION

- [Syntax | 1518](#)
- [Hierarchy Level | 1518](#)
- [Description | 1518](#)
- [Options | 1518](#)
- [Required Privilege Level | 1518](#)
- [Release Information | 1519](#)

Syntax

```
set paragon cluster papi papi-local-user-management (false|true)
```

Hierarchy Level

[edit]

Description

Enable or disable local user management on the nodes.

Options

papi-local-user-management (false true)	Set to false to disable local user management or set to true to enable local user management.

Required Privilege Level

configure

Release Information

Statement introduced in Juniper Paragon Automation Release 2.0.0.

RELATED DOCUMENTATION

[set paragon cluster mail-server | 1512](#)

[show configuration paragon cluster | 1439](#)

[delete paragon cluster | 1491](#)

set paragon monitoring

IN THIS SECTION

- [Syntax | 1519](#)
- [Hierarchy Level | 1519](#)
- [Description | 1520](#)
- [Options | 1520](#)
- [Required Privilege Level | 1522](#)
- [Release Information | 1522](#)

Syntax

```
set paragon monitoring sink id sink-options
```

```
set paragon monitoring source id source-options
```

Hierarchy Level

[edit]

Description

Configure the monitoring in Routing Director to collect metrics from different types of sources (host metrics, kube-state-metric, syslogs and so on) and forward the collected data to designated sinks, or destinations (Prometheus, VictoriaMetrics, and so on).

Use the request `paragon deploy monitoring` command to deploy the monitoring configuration. For more information, see *Post Installation Tasks*.

Options

<code>sink id elasticsearch elasticsearch-options</code>	For the specified sink ID, write the log messages to elasticsearch. The sink ID must be in the <code>[a-z][a-z0-9_-]*</code> format. <i>elasticsearch-options</i> include the following: • <code>api_version version</code>—API version of elasticsearch. Options include <code>auto</code>, <code>v6</code>, <code>v7</code>, and <code>v8</code>. • <code>compression method</code>—Data compression method. Default is <code>none</code>. Options include <code>none</code>, <code>gzip</code>, <code>snappy</code>, <code>zlib</code>, and <code>zstd</code>. • <code>endpoints (value / [values])</code>—HTTP(S) endpoint of sources/sinks. Enter a single <i>value</i> or multiple <i>[values]</i>. • <code>healthcheck (none enabled)</code>—When enabled, check the health of the sink when the Vector starts up. • <code>mode (bulk data_stream)</code>—Elasticsearch indexing mode. • <code>query parameters</code>—Custom parameters to add to the query string for each HTTP request sent to Elasticsearch, in the <code>arg1_key#arg1_value#arg2_key#arg2_value...</code> format. Number of hashtag-separated items has to be an even number.

sink <i>id</i> prometheus_remote_write <i>prometheus-options</i>	For the specified sink ID, write metric data to the endpoint with prometheus remote write protocol. The sink ID must be in the [a-z][a-z0-9_-]* format. <i>prometheus-options</i> include the following: • compression <i>method</i>—Data compression method. Default is snappy. Options include gzip, snappy, and zstd. • endpoint <i>value</i>—HTTP(S) endpoint of sources/sinks. • healthcheck (none enabled)—When enabled, check the health of the sink when the Vector starts up.
source <i>id</i> cluster kafka <i>options</i>	For the specified source ID, monitor cluster-level metric source and subscribe to kafka topics for messages. The default sources are audit, host, and ksm. <i>options</i> include the following options: • bootstrap_servers <i>servers</i>—Comma separated list of bootstrap servers. • group_id <i>value</i>— Kafka subscription group ID. • topics ((<i>value</i> / [<i>values</i>] / audits-dev)—List of kafka topics. Enter a single <i>value</i> or multiple [<i>values</i>] or audits-dev.
source <i>id</i> cluster prometheus_scrape <i>options</i>	For the specified source ID, monitor cluster-level metric source and collect metrics through sources which writes prometheus-format data. The default sources are audit, host, and ksm. <i>options</i> include the following: • endpoints (<i>value</i> / [<i>values</i>])—HTTP(S) endpoint of sources/sinks. Enter a single <i>value</i> or multiple [<i>values</i>]. • scrape_interval_secs <i>interval</i>— Scrape interval in seconds.

source *id* node host_metrics scrape_interval_secs
interval

For the specified source ID, monitor and collect node-level metric data. Enter scrape interval in seconds. The default sources are audit, host, and ksm.

Required Privilege Level

configure

Release Information

Statement introduced in Juniper Paragon Automation Release 2.0.0.

RELATED DOCUMENTATION

[request paragon deploy](#) | [1400](#)

set system login

IN THIS SECTION

- [Syntax](#) | [1522](#)
- [Hierarchy Level](#) | [1523](#)
- [Description](#) | [1523](#)
- [Options](#) | [1523](#)
- [Required Privilege Level](#) | [1523](#)
- [Release Information](#) | [1524](#)

Syntax

```
set system login user username class (read-only|super-user) authentication authentication_method
```

Hierarchy Level

[edit]

Description

Set the names, login classes and passwords for users.

Options

user <i>username</i>	Set the login privileges for the specified username.
class (read-only super-user)	Set the access privilege for the specified user.
authentication <i>authentication_method</i>	Set the authentication methods that can be used to authenticate user. Options are: • plain-text-password—Enter the password as plain text when prompted. • encrypted-password "encrypted_password_string"—Enter an encrypted password string in quotation marks. The supported encryption algorithm is SHA-512. • ssh-ecdsa <i>authorized-key-ecdsa</i>—Enter the SSH ECDSA public key string. • ssh-ed25519 <i>authorized-key-ed25519</i>—Enter the SSH ED25519 public key string. • ssh-rsa <i>authorized-key-rsa</i>—Enter the SSH RSA public key string.

Required Privilege Level

configure

Release Information

Statement introduced in Juniper Paragon Automation Release 2.0.0.

show paragon cluster (c)

IN THIS SECTION

- [Syntax | 1525](#)
- [Hierarchy Level | 1525](#)
- [Description | 1525](#)
- [Options | 1526](#)
- [Required Privilege Level | 1529](#)
- [Release Information | 1529](#)

Syntax

```
show paragon cluster
```

```
show paragon cluster applications (none|applications_options)
```

```
show paragon cluster common-services (none|common_services_options)
```

```
show paragon cluster insights victoria-metrics (none|victoria_metrics_options)
```

```
show paragon cluster install (none|offline-install|prefer-ipv6)
```

```
show paragon cluster kubernetes (none|address-family [none|cluster-ipv6-enabled])
```

```
show paragon cluster mail-server (none|mail_server_options)
```

```
show paragon cluster nodes (none|nodes_options)
```

```
show paragon cluster ntp (none|ntp-servers)
```

```
show paragon cluster papi (none|papi-local-user-management)
```

Hierarchy Level

[edit]

Description

Display all the configured parameters under the paragon hierarchy.

Options

none	Show all the configured parameters under the paragon hierarchy.
applications (none <i>applications_options</i>)	Show Paragon Active Assurance Test Agent gateway virtual IP address, hostname, and default Web administrator login credentials. <i>applications_options</i> options include: • active-assurance • active-assurance test-agent-gateway-vip • active-assurance test-agent-gateway-vip-ipv6 • active-assurance test-agent-gateway-hostname • web-ui • web-ui web-admin-password • web-ui web-admin-user • pathfinder pce-server pce-server-vip For more information, see "set paragon cluster applications" on page 1498 and "show paragon cluster applications (c)" on page 1529.

common-services (none <i>common_services_options</i>)	Show the common ingress virtual IP (VIP) address, hostname, and custom user certificates usage. <i>common_services_options</i> options include: • ingress ingress-hostname • ingress ingress-vip • ingress ingress-vip-ipv6 • ingress user-certificate • ingress user-certificate use-user-certificate • ingress user-certificate user-certificate-filename • ingress user-certificate-key-filename For more information, see "set paragon cluster common-services" on page 1504 and "show paragon cluster common-services (c)" on page 1532.
insights (none victoria-metrics <i>victoria_metrics_options</i>)	Display Victoria Metrics configuration parameters. <i>victoria_metrics_options</i> options include: • vm-cluster-replication-factor • vm-cluster-retention-period • vm-cluster-storage-instances For more information, see "set paragon cluster insights victoria-metrics" on page 1507.
install (none offline-install prefer-ipv6)	Display the cluster installation parameters and whether the Routing Director cluster is enabled to be installed in an air-gap environment. For more information, see "set paragon cluster install" on page 1509.

kubernetes (none address-family [none cluster-ipv6-enabled])	Display whether the usage of IPv6 addresses is enabled or disabled for the cluster. For more information, see "set paragon cluster kubernetes" on page 1511.
mail-server (none <i>mail_server_options</i>)	Display SMTP configuration settings. <i>mail_server_options</i> options include: • smtp-allowed-sender-domains • smtp-relayhost • smtp-relayhost-password • smtp-relayhost-username • smtp-sender-email • smtp-sender-name For more information, see "set paragon cluster mail-server" on page 1512.
nodes (none <i>nodes_options</i>)	Display IP address or hostname of the Routing Director cluster node and the index number of the node. <i>nodes_options</i> options include: • kubernetes <i>index</i> • kubernetes <i>index</i> address • kubernetes <i>index</i> hostname For more information, see "set paragon cluster nodes" on page 1515.
ntp (none ntp-servers)	Display the configured NTP servers. For more information, see "set paragon cluster ntp" on page 1516.

papi (none papi-local-user-management)	Display whether local user management on the nodes is enabled or disabled. For more information, see "set paragon cluster papi" on page 1518.

Required Privilege Level

configure

Release Information

- Statement introduced in Juniper Paragon Automation Release 2.0.0.
- pathfinder option introduced in Juniper Paragon Automation Release 2.1.0.
- insights option introduced in Juniper Paragon Automation Release 2.1.0.
- kubernetes option introduced in Juniper Paragon Automation Release 2.2.0.
- prefer-ipv6 option introduced in Juniper Paragon Automation Release 2.2.0.
- ingress ingress-vip-ipv6 option introduced in Juniper Paragon Automation Release 2.2.0.
- test-agent-gateway-vip-ipv6 option introduced in Juniper Paragon Automation Release 2.2.0.

RELATED DOCUMENTATION

delete paragon cluster 1491
show paragon cluster applications (c) 1529
show paragon cluster common-services (c) 1532

show paragon cluster applications (c)

IN THIS SECTION

- Syntax | [1530](#)

- [Hierarchy Level | 1531](#)
- [Description | 1531](#)
- [Options | 1531](#)
- [Required Privilege Level | 1532](#)
- [Release Information | 1532](#)

Syntax

```
show paragon cluster applications
```

```
show paragon cluster applications active-assurance (none|test-agent-gateway-vip|test-agent-gateway-vip-ipv6|test-agent-gateway-hostname)
```

```
show paragon cluster applications aiops (none|enable-device-health|install-aiml)
```

```
show paragon cluster applications gnmi-term (none|gnmi-term-host|gnmi-term-host-ipv6)
```

```
show paragon cluster applications oc-term (none|oc-term-host|oc-term-host-ipv6|oc-term-previous-management-ip-family)
```

```
show paragon cluster applications pathfinder pce-server pce-server-vip
```

```
show paragon cluster applications routingbot (none|install-routingbot|routingbot-crpd-vip)
```

```
show paragon cluster applications web-ui (none|web-admin-password|web-admin-user)
```

Hierarchy Level

[edit]

Description

Show all configured applications details.

Options

none	Show all configured applications details.
active-assurance (none test-agent-gateway-vip test-agent-gateway-vip-ipv6 test-agent-gateway-hostname)	Show all the configured Paragon Active Assurance Test Agent gateway details, such as IPv4 and IPv6 addresses, and hostname.
aiops (none enable-device-health install-aiml)	Show AI/ML (artificial intelligence [AI] and machine learning [ML]) configuration settings including device-health monitoring settings.
gnmi-term (none gnmi-term-host gnmi-term-host-ipv6)	Show the configured IPv4 or IPv6 VIP addresses for gNMI.
oc-term (none oc-term-host oc-term-host-ipv6 oc-term-previous-management-ip-family)	Show the configured IPv4 or IPv6 VIP addresses for OC-TERM.
pathfinder pce-server pce-server-vip	Show the configured PCE server VIP address details.
routingbot (none install-routingbot routingbot-crpd-vip)	Show routing observability feature configuration settings including the CRPD VIP address.
web-ui (none web-admin-password web-admin-user)	Show the Web administrator login credentials. The password is encrypted. The Web administrator is the first user to log in to the Paragon Automation GUI and has super user privileges.

Required Privilege Level

configure

Release Information

Statement introduced in Juniper Paragon Automation Release 2.0.0.

test-agent-gateway-vip-ipv6 option introduced in Juniper Paragon Automation Release 2.2.0.

pathfinder option introduced in Juniper Paragon Automation Release 2.1.0.

aiops, gnmi-term, oc-term, and routingbot options introduced in Juniper Paragon Automation Release 2.4.0.

show paragon cluster common-services (c)

IN THIS SECTION

- [Syntax | 1533](#)
- [Hierarchy Level | 1533](#)
- [Description | 1533](#)
- [Options | 1533](#)
- [Required Privilege Level | 1535](#)
- [Release Information | 1535](#)

Syntax

```
show paragon cluster common-services
```

```
show paragon cluster common-services ingress (none|ingress-hostname|ingress-vip|ingress-vip-ipv6|  
user-certificate)
```

```
show paragon cluster common-services ingress user-certificate use-user-certificate (none|true|  
false)
```

```
show paragon cluster common-services ingress user-certificate user-certificate-filename
```

```
show paragon cluster common-services ingress user-certificate user-certificate-key-filename
```

```
show paragon cluster common-services metallb metallb-bgp-peer peer-ip v4-address (none|peer-asn|  
local-asn|local-nodes)
```

```
show paragon cluster common-services metallb metallb-bgp-peer-ipv6 peer-ip v6 address (none|peer-  
asn|local-asn|local-nodes)
```

Hierarchy Level

[edit]

Description

Show all common ingress configuration information.

Options

(none ingress)	Show all common ingress configuration information.
ingress ingress-hostname	Show the hostname for the ingress VIP address.
ingress ingress-vip	Show the common ingress IPv4 VIP address.
ingress ingress-vip-ipv6	Show the common ingress IPv6 VIP address.
ingress user-certificate	Show custom user certificates configuration.
ingress user-certificate use-user-certificate (none true false)	Show custom user certificates configuration.
ingress user-certificate user-certificate-filename	Show the custom user certificate filename.
ingress user-certificate-key-filename	Show the custom user certificate key filename.
metallb metallb-bgp-peer peer-ip <i>v4-address</i>	Show all the configured BGP peering information between the specified top-of-rack (ToR) router and the cluster nodes connected to it. Here, <i>v4-address</i> is the IP address of the ToR router.
metallb metallb-bgp-peer peer-ip <i>v4-address</i> peer-asn	Show the configured AS number for the specified ToR router.
metallb metallb-bgp-peer peer-ip <i>v4-address</i> local-asn	Show the configured AS number for the cluster nodes connected to the specified ToR router.
metallb metallb-bgp-peer peer-ip <i>v4-address</i> local-nodes	Show the cluster nodes connected to the specified ToR router.
metallb metallb-bgp-peer-ipv6 peer-ip <i>v6 address</i>	Show all the configured BGP peering information between the specified top-of-rack (ToR) router and the cluster nodes connected to it. Here, <i>v6-address</i> is the IPv6 address of the ToR router.

metallb metallb-bgp-peer-ipv6 peer-ip <i>v6 address</i> peer-asn	Show the configured AS number for the specified ToR router.
metallb metallb-bgp-peer-ipv6 peer-ip <i>v6 address</i> local-asn	Show the configured AS number for the cluster nodes connected to the specified ToR router.
metallb metallb-bgp-peer-ipv6 peer-ip <i>v6 address</i> local-nodes	Show the cluster nodes connected to the specified ToR router.

Required Privilege Level

configure

Release Information

Statement introduced in Juniper Paragon Automation Release 2.0.0.

ingress-vip-ipv6 option introduced in Juniper Paragon Automation Release 2.2.0.

metallb option added in Juniper Paragon Automation Release 2.4.0.

Troubleshooting Commands

IN THIS CHAPTER

- [Troubleshoot Using the Paragon Shell CLI Commands | 1536](#)
- [request paragon support information | 1543](#)
- [request paragon troubleshooting information | 1546](#)
- [request paragon debug | 1550](#)
- [request paragon debug get-tsdb-data | 1552](#)
- [request paragon debug insights-kafka-data | 1554](#)
- [request paragon debug kafka | 1556](#)
- [request paragon debug logs | 1558](#)
- [request paragon debug logs namespace | 1561](#)
- [request paragon debug postgres | 1571](#)
- [request paragon debug redis | 1573](#)

Troubleshoot Using the Paragon Shell CLI Commands

SUMMARY

Read the following topics to understand how to use the Paragon Shell CLI commands to troubleshoot issues with your Routing Director cluster.

IN THIS SECTION

- [Overview | 1537](#)
- [request paragon support information | 1537](#)
- [request paragon troubleshooting information | 1538](#)
- [Other Troubleshooting Commands to Debug Issues | 1541](#)

Overview

Routing Director enables you to troubleshoot Routing Director cluster-related issues using Paragon Shell CLI commands. To execute these Paragon Shell CLI troubleshooting commands, you must log in to any of the cluster nodes. When you execute these commands, a series of commands are executed one after the other.

The Paragon Shell CLI troubleshooting commands are:

- ["request paragon support information" on page 1537](#)
- ["request paragon troubleshooting information" on page 1538](#)

Routing Director also enables you to debug issues by collecting data from the Redis database, Kafka messages, service logs, and the time series database (TSDB). These commands are not part of the Paragon Shell CLI troubleshooting commands, and must be run separately. See ["Other Troubleshooting Commands to Debug Issues" on page 1541](#) for more information.

request paragon support information

The request paragon support information command displays an in-depth status report of your Routing Director cluster configuration.

The show commands that are executed when you run the request paragon support information command are listed in [Table 350 on page 1537](#).

Table 350: request paragon support information Commands

Command	Description
<code>show paragon cluster nodes</code>	Show node information of your Routing Director cluster.
<code>show paragon cluster pods</code>	Show pod information of your Routing Director cluster.
<code>show paragon cluster namespaces</code>	Show namespace information of your Routing Director cluster.
<code>show paragon cluster details</code>	Show storage and controller node information of your Routing Director cluster.

Table 350: request paragon support information Commands (Continued)

Command	Description
<code>show paragon version</code>	Show the version of your Routing Director cluster.
<code>show paragon images version</code>	Show the version of pods in your Routing Director cluster.
<code>show paragon cluster pods namespace healthbot sort memory</code>	Show the top pods of the healthbot namespace sorted by memory utilization.
<code>show paragon cluster pods namespace healthbot sort cpu</code>	Show the top pods of the healthbot namespace sorted by CPU utilization.
<code>show paragon pvc details</code>	Show the persistent volume (PV) and persistent volume claim (PVC) information.

The `request paragon support information` command also runs many `kubectl` commands. These commands provide you debugging information such as Helm deployment service information for the NorthStar namespace, Kubernetes deployment information for api-aggregator service and so on.

request paragon troubleshooting information

The `request paragon troubleshooting information` command provides troubleshooting information of the `ems`, `foghorn`, `insights`, `paa`, `trust`, and `pathfinder` Routing Director services.

To view the list of available services, run the following command:

```
user@node> request paragon troubleshooting information service ?
```

Possible completions:

ems	ems service
foghorn	foghorn service
insights	insights service
paa	paa service
pathfinder	pathfinder service
trust	trust service

When you run the request paragon troubleshooting information command, a troubleshooting_<date_time>.tar.gz file is generated. You can share this file with the [Juniper Technical Assistance Center \(JTAC\)](#) for further evaluation. This .tar.gz file is saved in the /root/troubleshooting/ directory.

The commands that are executed when you run the request paragon troubleshooting information command are listed in [Table 351 on page 1539](#).

Table 351: request paragon troubleshooting information Commands

Command	Description
"request paragon debug logs namespace healthbot service tsdb-shim" on page 1561	Generate a log file of the tsdb-shim service within the healthbot namespace.
"request paragon debug logs namespace healthbot service tand" on page 1561	Generate a log file of the tand service within the healthbot namespace.
"request paragon debug logs namespace healthbot service jtimon" on page 1561	Generate a log file of the jtimon service within the healthbot namespace.
"request paragon debug logs namespace healthbot service config-server container config-server" on page 1561	Generate a log file of the config-server service within the healthbot namespace.
"request paragon debug logs namespace healthbot service api-server" on page 1561	Generate a log file of the api-server service within the healthbot namespace.
"request paragon debug logs namespace healthbot service analytical-engine" on page 1561	Generate a log file of the analytical-engine service within the healthbot namespace.
"request paragon debug logs namespace healthbot service alerta" on page 1561	Generate a log file of the alerta service within the healthbot namespace.
"request paragon debug postgres database postgres username healthbot measurement db_counters output file" on page 1561	Generate a text file (JSON format) with Postgres information.
"request paragon debug logs namespace foghorn service order-management" on page 1563	Generate a log file of the order-management service within the foghorn namespace.

Table 351: request paragon troubleshooting information Commands *(Continued)*

Command	Description
"request paragon debug logs namespace foghorn service placement" on page 1563	Generate a log file of the placement service within the foghorn namespace.
"request paragon debug logs namespace foghorn service cmgd" on page 1563	Generate a log file of the cmgd service within the foghorn namespace.
"request paragon debug logs namespace airflow service workflow-manager" on page 1569	Generate a log file of the workflow-manager service within the airflow namespace.
"request paragon debug logs namespace northstar service toposerver" on page 1567	Generate a log file of the toposerver service within the northstar namespace.
"request paragon debug logs namespace northstar service configmonitor" on page 1567	Generate a log file of the configmonitor service within the northstar namespace.
"request paragon debug logs namespace northstar service web" on page 1567	Generate a log file of the web service within the northstar namespace.
"request paragon debug logs namespace northstar service api-aggregator" on page 1567	Generate a log file of the api-aggregator service within the northstar namespace.
"request paragon debug logs namespace papi service oc-term" on page 1565	Generate a log file of the oc-term service within the papi namespace.
"request paragon debug logs namespace papi service papi" on page 1565	Generate a log file of the papi service within the papi namespace.
"request paragon debug logs namespace papi service papi-ws" on page 1565	Generate a log file of the papi-ws service within the papi namespace.

Other Troubleshooting Commands to Debug Issues

You can run commands to collect data from the Redis database, Kafka messages, service logs, and the time series database (TSDB). You can use this data to troubleshoot issues with your Routing Director cluster. These commands are not part of the Paragon Shell CLI troubleshooting commands, and must be run separately. [Table 352 on page 1541](#) lists the commands.

Table 352: Alternative Commands to Debug Issues

Command	Description
<i>Kafka</i>	
request paragon debug kafka ?	Display possible completions for the request paragon debug kafka command.
"request paragon debug kafka options "-C -t <i>topic-name</i> -o s@ <i>start-time</i> -o e@ <i>end-time</i> -e -JB" output-file " <i>file-name</i> " on page 1556	Generate an output file of Kafka messages for a topic for a specified period of time.
<i>Insights Kafka</i>	
request paragon debug insights-kafka-data ?	Display possible completions for the request paragon debug insights-kafka-data command.
"request paragon debug insights-kafka-data device " <i>device-id</i> " time-period " <i>duration</i> " on page 1554	Display insights-kafka-data information for a device, for a specific time period. An output file of the information is generated.
<i>Redis</i>	
request paragon debug redis ?	Display possible completions for the request paragon debug redis command.
"request paragon debug redis redis-key-pattern "insights" on page 1573	Display Redis key pattern information for redis-keys with pattern "insights".

Table 352: Alternative Commands to Debug Issues *(Continued)*

Command	Description
"request paragon debug redis-key-pattern "insights" output file" on page 1573	Generate an output file of Redis key pattern information for redis-keys with pattern "insights".
<i>Service Logs</i>	
request paragon debug logs ?	Display possible completions for the request paragon debug logs command.
"request paragon debug logs namespace <i>name</i> service <i>service-name</i> time <i>duration</i> " on page 1558	Generate a log file for a service within a namespace for the specified time period.
<i>TSDB</i>	
request paragon debug get-tsdb-data ?	Display possible completions for the request paragon debug get-tsdb-data command.
"request paragon debug get-tsdb-data device <i>device-id</i> topic "topic-name" output <i>file</i> " on page 1552	Generate an output file of TSDB data for a particular device.
<i>Postgres</i>	
request paragon debug postgres ?	Display possible completions for the request paragon debug postgres command.
"request paragon debug postgres database <i>database-name</i> username <i>username</i> measurement <i>measurement-name</i> output (file)" on page 1571	Generate an output file of the measurement value information of the Postgres database.

request paragon support information

IN THIS SECTION

- [Syntax | 1543](#)
- [Description | 1543](#)
- [Options | 1544](#)
- [Required Privilege Level | 1545](#)
- [Output Fields | 1545](#)
- [Sample Output | 1545](#)
- [Release Information | 1546](#)

Syntax

```
request paragon support information
```

Description

Show an in-depth status report of your Paragon Automation cluster configuration.

When you enter the request paragon support information command, the following show commands are executed.

Table 353: request paragon support information Commands

Command	Description
<code>show paragon cluster nodes</code>	Show node information of your Routing Director cluster.
<code>show paragon cluster pods</code>	Show pod information of your Routing Director cluster.

Table 353: request paragon support information Commands (Continued)

Command	Description
<code>show paragon cluster namespaces</code>	Show namespace information of your Routing Director cluster.
<code>show paragon cluster details</code>	Show storage and controller node information of your Routing Director cluster.
<code>show paragon version</code>	Show the version of your Routing Director cluster.
<code>show paragon images version</code>	Show the version of pods in your Routing Director cluster.
<code>show paragon cluster pods namespace healthbot sort memory</code>	Show the top pods of the healthbot namespace sorted by memory utilization.
<code>show paragon cluster pods namespace healthbot sort cpu</code>	Show the top pods of the healthbot namespace sorted by CPU utilization.
<code>show paragon pvc details</code>	Show the persistent volume (PV) and persistent volume claim (PVC) information.

The `request paragon support information` command also runs many `kubectl` commands. These commands provide you debugging information such as Helm deployment service information for the NorthStar namespace, Kubernetes deployment information for api-aggregator service and so on.

To archive the status report of your cluster configuration, you can use the `request paragon support information | save paragon/troubleshooting/rsi.txt` command. The `rsi.txt` file is located under the `/root/troubleshooting` directory of the Linux root shell. Type `exit` to exit to the Linux root shell.

Options

<code>none</code>	Show in-depth status report of your cluster configuration.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided with detailed status information of your cluster.

Sample Output

request paragon support information

```
user@node> request paragon support information

user@node> show paragon cluster nodes

NAME      STATUS  ROLES                AGE  VERSION
eop-250   Ready   control-plane,etcd,master 36h  v1.28.6+rke2r1
eop-251   Ready   control-plane,etcd,master 36h  v1.28.6+rke2r1
eop-252   Ready   control-plane,etcd,master 36h  v1.28.6+rke2r1
eop-253   Ready   influxdb-worker,worker  36h  v1.28.6+rke2r1

user@node> show paragon cluster pods

NAMESPACE      NAME                                     READY
STATUS  RESTARTS      AGE
airflow      airflow-scheduler-9d86ffb65-gdn7h      2/2
Running    3 (34h ago)   36h
airflow      airflow-webserver-fdbdcf8-2plkj        1/1
Running    0              36h
airflow      airflow-worker-6c568dc6c8-n541x        1/1
Running    1 (36h ago)   36h
airflow      airflow-worker-6c568dc6c8-qlw9z        1/1
Running    0              36h
airflow      workflow-manager-7574d7cc95-6ps66      1/1
Running    0              36h
arango       arango-arango-operator-58dbbf6669-9hn92 1/1
Running    1 (36h ago)   36h
arango       arango-arango-operator-58dbbf6669-q74sg 1/1
Running    3 (34h ago)   36h
arango       foghorn-agnt-8xpiqppm-f450c7
```

<Output Snipped>

Release Information

Command introduced in Juniper Paragon Automation Release 2.1.0.

RELATED DOCUMENTATION

[Troubleshoot Using the Paragon Shell CLI Commands](#) | [1536](#)

request paragon troubleshooting information

IN THIS SECTION

- [Syntax](#) | [1546](#)
- [Description](#) | [1547](#)
- [Options](#) | [1549](#)
- [Required Privilege Level](#) | [1549](#)
- [Output Fields](#) | [1549](#)
- [Sample Output](#) | [1549](#)
- [Release Information](#) | [1550](#)

Syntax

```
request paragon troubleshooting information
```

```
request paragon troubleshooting information service name
```

Description

Generate a log file of various Paragon Automation services. You can share this file with the [Juniper Technical Assistance Center \(JTAC\)](#) for further evaluation.

When you run the request paragon troubleshooting information command, the following commands are executed.

Table 354: request paragon troubleshooting information Commands

Command	Description
"request paragon debug logs namespace healthbot service tsdb-shim" on page 1561	Generate a log file of the tsdb-shim service within the healthbot namespace.
"request paragon debug logs namespace healthbot service tand" on page 1561	Generate a log file of the tand service within the healthbot namespace.
"request paragon debug logs namespace healthbot service jtimon" on page 1561	Generate a log file of the jtimon service within the healthbot namespace.
"request paragon debug logs namespace healthbot service config-server container config-server" on page 1561	Generate a log file of the config-server service within the healthbot namespace.
"request paragon debug logs namespace healthbot service api-server" on page 1561	Generate a log file of the api-server service within the healthbot namespace.
"request paragon debug logs namespace healthbot service analytical-engine" on page 1561	Generate a log file of the analytical-engine service within the healthbot namespace.
"request paragon debug logs namespace healthbot service alerta" on page 1561	Generate a log file of the alerta service within the healthbot namespace.
"request paragon debug postgres database postgres username healthbot measurement db_counters output file" on page 1561	Generate a text file (JSON format) with Postgres information.

Table 354: request paragon troubleshooting information Commands (Continued)

Command	Description
"request paragon debug logs namespace foghorn service order-management" on page 1563	Generate a log file of the order-management service within the foghorn namespace.
"request paragon debug logs namespace foghorn service placement" on page 1563	Generate a log file of the placement service within the foghorn namespace.
"request paragon debug logs namespace foghorn service cmgd" on page 1563	Generate a log file of the cmgd service within the foghorn namespace.
"request paragon debug logs namespace airflow service workflow-manager" on page 1569	Generate a log file of the workflow-manager service within the airflow namespace.
"request paragon debug logs namespace northstar service toposerver" on page 1567	Generate a log file of the toposerver service within the northstar namespace.
"request paragon debug logs namespace northstar service configmonitor" on page 1567	Generate a log file of the configmonitor service within the northstar namespace.
"request paragon debug logs namespace northstar service web" on page 1567	Generate a log file of the web service within the northstar namespace.
"request paragon debug logs namespace northstar service api-aggregator" on page 1567	Generate a log file of the api-aggregator service within the northstar namespace.
"request paragon debug logs namespace papi service oc-term" on page 1565	Generate a log file of the oc-term service within the papi namespace.
"request paragon debug logs namespace papi service papi" on page 1565	Generate a log file of the papi service within the papi namespace.
"request paragon debug logs namespace papi service papi-ws" on page 1565	Generate a log file of the papi-ws service within the papi namespace.

A `troubleshooting_date_time.tar.gz` file for the various Paragon Automation services is also created. This `.tar.gz` file is saved in the `/root/troubleshooting/` directory.

Options

<code>none</code>	Generate log files for all Paragon Automation services.
<code>service name</code>	Generate log file for a service.

Required Privilege Level

view

Output Fields

When you run this command, a log file is generated. You can use the log file to troubleshoot issues with your Paragon Automation cluster.

Sample Output

request paragon troubleshooting information

```
user@node> request paragon troubleshooting information

user@node> request paragon debug logs namespace healthbot service tsdb-shim app insights
Log file for tsdb-shim created at: /root/troubleshooting/app/insights/logs/logs-tsdb-
shim-2024-03-03_14-05-21.tar.gz

user@node> request paragon debug logs namespace healthbot service tand app insights
Log file for tand created at: /root/troubleshooting/app/insights/logs/logs-
tand-2024-03-03_14-05-24.tar.gz

user@node> request paragon debug logs namespace healthbot service jtimon app insights
Log file for jtimon created at: /root/troubleshooting/app/insights/logs/logs-
jtimon-2024-03-03_14-05-26.tar.gz

user@node> request paragon debug postgres database postgres username healthbot measurement
db_counters output file app insights
```

```
Writing postgres output at: /root/troubleshooting/app/insights/postgres/postgres_2024-03-03_14-05-43.txt
```

```
user@node> request paragon debug logs namespace foghorn service order-management app foghorn
Log file for order-management created at: /root/troubleshooting/app/foghorn/logs/logs-order-management-2024-03-03_14-05-46.tar.gz
```

```
user@node> request paragon debug logs namespace foghorn service placement app foghorn
Log file for placement created at: /root/troubleshooting/app/foghorn/logs/logs-placement-2024-03-03_14-05-50.tar.gz
```

<Output Snipped>

Release Information

Command introduced in Juniper Paragon Automation Release 2.0.0.

RELATED DOCUMENTATION

[Troubleshoot Using the Paragon Shell CLI Commands](#) | 1536

request paragon debug

IN THIS SECTION

- [Syntax](#) | 1551
- [Description](#) | 1551
- [Options](#) | 1551
- [Required Privilege Level](#) | 1552
- [Output Fields](#) | 1552
- [Release Information](#) | 1552

Syntax

```
request paragon debug (database / service / logs)
```

Description

Generate an output file for a database, service, or log.

Options

get-tsdb-data	Generate an output file of TSDB information. See "request paragon debug get-tsdb-data" on page 1552.
insights-kafka-data	Generate an output file of insights-kafka-data information for a device for a specified period of time. See "request paragon debug insights-kafka-data" on page 1554.
kafka	Generate an output file of Kafka messages for a topic for a specified period of time. See "request paragon debug kafka " on page 1556.
logs	Generate a log file for a service within a namespace for a specified period of time. See "request paragon debug logs" on page 1558.
postgres	Generate an output file of the measurement value information of the Postgres database. See "request paragon debug postgres" on page 1571.
redis	Generate an output file of Redis key pattern information for redis-keys. See "request paragon debug redis" on page 1573.

Required Privilege Level

view

Output Fields

When you enter this command, an output file is generated for the database, service, or log.

Release Information

Command introduced in Juniper Paragon Automation Release 2.0.0.

request paragon debug get-tsdb-data

IN THIS SECTION

- [Syntax | 1552](#)
- [Description | 1553](#)
- [Options | 1553](#)
- [Output Fields | 1553](#)
- [Sample Output | 1553](#)
- [Required Privilege Level | 1554](#)
- [Release Information | 1554](#)

Syntax

```
request paragon debug get-tsdb-data device device-id topic "topic-name" output file
```

```
request paragon debug get-tsdb-data device device-id topic "topic-name" output file last n
```

Description

Generate time series database (TSDB) information and export data to an output file.

Options

device <i>device-id</i>	(Mandatory) Device Id of the device.
topic " <i>topic-name</i> "	(Mandatory) Name of the topic for which the TSDB data is exported.
last <i>n</i>	(Optional) Export last <i>n</i> seconds of data.
output <i>file</i>	(Mandatory) Redirect output to file.

Output Fields

When you enter this command, a TSDB data for a topic is generated and exported to an output file.

Sample Output

request paragon debug get-tsdb-data device 2c6bf5f19f00 topic "routing.fib" output file

```
user@node> request paragon debug get-tsdb-data device 2c6bf5f19f00 topic "routing.fib" output
file
Writing tsdb output at: /root/troubleshooting/tsdb/tsdb_2024-04-17_01-58-02.txt
```

request paragon debug get-tsdb-data device 2c6bf5f19f00 topic "routing.fib" output file last 120

```
user@node> request paragon debug get-tsdb-data device 2c6bf5f19f00 topic "routing.fib" output
file last 120
Writing tsdb output at: /root/troubleshooting/tsdb/tsdb_2024-04-17_01-58-02.txt
```

Required Privilege Level

view

Release Information

Command introduced in Juniper Paragon Automation Release 2.0.0.

request paragon debug insights-kafka-data

IN THIS SECTION

- [Syntax | 1554](#)
- [Description | 1554](#)
- [Options | 1554](#)
- [Required Privilege Level | 1555](#)
- [Output Fields | 1555](#)
- [Sample Output | 1555](#)
- [Release Information | 1556](#)

Syntax

```
request paragon debug insights-kafka-data device "device-id" time-period "duration"
```

Description

Show insights-kafka-data information for a device and for a specified period of time. An output file of the information is generated.

Options

device " <i>device-id</i> "	Device ID of the device.
time-period " <i>duration</i> "	Time period (in minutes) for which insights-kafka-data needs to run. Example: 4m

Required Privilege Level

view

Output Fields

When you enter this command, insights-kafka-data messages for a device and for a specified time period is displayed. The location and name of the output file is also displayed.

Sample Output

request paragon debug insights-kafka-data device "111111" time-period "1m"

```
user@node> request paragon debug insights-kafka-data device "111111" time-period "1m"
kubectl --kubeconfig /etc/kubernetes/admin.conf exec common-utils-7c995f68dc-4vnhf -n common
-- /bin/bash -c 'cd /common-utils && ./insights-data-listener -brokers kafka.common:9092 -topics
gnmi-data-dev,insights-ingest-data-dev,insights-topic-rule-data-dev -device 111111'
/usr/sbin/cli -c "ssh node commands-to-execute \"kubectl --kubeconfig /etc/kubernetes/admin.conf
cp common-utils-7c995f68dc-4vnhf:/common-utils/insights-ingest-data-dev /root/troubleshooting/
insights_kafka/insights-ingest-data-dev_2024-04-17_02-47-22 -n common\""
tar: Removing leading '/' from member names
Writing kafka output at: /root/troubleshooting/insights_kafka/insights-ingest-data-
dev_2024-04-17_02-47-22
/usr/sbin/cli -c "ssh node commands-to-execute \"kubectl --kubeconfig /etc/kubernetes/admin.conf
cp common-utils-7c995f68dc-4vnhf:/common-utils/insights-topic-rule-data-dev /root/
troubleshooting/insights_kafka/insights-topic-rule-data-dev_2024-04-17_02-47-22 -n common\""
tar: Removing leading '/' from member names
Writing kafka output at: /root/troubleshooting/insights_kafka/insights-topic-rule-data-
dev_2024-04-17_02-47-22
/usr/sbin/cli -c "ssh node commands-to-execute \"kubectl --kubeconfig /etc/kubernetes/admin.conf
cp common-utils-7c995f68dc-4vnhf:/common-utils/gnmi-data-dev /root/troubleshooting/
insights_kafka/gnmi-data-dev_2024-04-17_02-47-22 -n common\""
```

```
tar: Removing leading '/' from member names
Writing kafka output at: /root/troubleshooting/insights_kafka/gnmi-data-dev_2024-04-17_02-47-22
```

Release Information

Command introduced in Juniper Paragon Automation Release 2.0.0.

request paragon debug kafka

IN THIS SECTION

- [Syntax | 1556](#)
- [Description | 1556](#)
- [Options | 1556](#)
- [Required Privilege Level | 1557](#)
- [Output Fields | 1557](#)
- [Sample Output | 1557](#)
- [Release Information | 1558](#)

Syntax

```
request paragon debug kafka options "query" output-file "file-name"
```

```
request paragon debug kafka options "query" output-file "file-name" time period duration
```

Description

Generate an output file of the Kafka messages for a topic, and for a specified period of time.

Options

options " <i>query</i> "	Add a Kafka query. Format: -C -t <i>topic-name</i> -o s@<i>start-time</i> -o e@<i>end-time</i> -e For example, to query a topic name from a particular start time, use -C -t netconf-resp-dev -o s@1713321079000 -e
output-file " <i>file-name</i> "	Name of the file where Kafka messages are stored.
time period <i>duration</i>	Duration for which kafkacat needs to run. Format: Specify integer followed by s/m/h/d/w/y (seconds/minutes/hours/days/weeks/years). For example, 2s.

Required Privilege Level

view

Output Fields

When you enter this command, an output file of the Kafka messages is created for the specified period of time.

Sample Output

request paragon debug kafka options "-C -t pf-cfg-topology -o s@1701705745000 -o e@1702569745000 -e -JB" output-file "file"

```
user@node> request paragon debug kafka options "-C -t pf-cfg-topology -o s@1701705745000 -o e@1702569745000 -e -JB" output-file "file"

Writing kafka output at: /root/troubleshooting/kafka/file
```

request paragon debug kafka options "-C -t netconf-resp-dev -o s@1713321079000 -e"
output-file test-netconf-resp-dev time-period 60s

```
user @node> request paragon debug kafka options "-C -t netconf-resp-dev -o s@1713321079000 -e"
output-file test-netconf-resp-dev time-period 60s
kubectl --kubeconfig /etc/kubernetes/admin.conf exec common-utils-dd49dd8d6-5ksks -n common
-- /bin/bash -c 'kcat -b $KAFKA_BROKER_ADDRESS -C -t netconf-resp-dev -o s@1713321079000 -e' > /
root/troubleshooting/kafka/test-netconf-resp-dev
Writing kafka output at: /root/troubleshooting/kafka/test-netconf-resp-dev
```

Release Information

Command introduced in Juniper Paragon Automation Release 2.0.0.

request paragon debug logs

IN THIS SECTION

- [Syntax | 1559](#)
- [Description | 1559](#)
- [Options | 1559](#)
- [Required Privilege Level | 1559](#)
- [Output Fields | 1560](#)
- [Sample Output | 1560](#)
- [Release Information | 1560](#)

Syntax

```
request paragon debug logs container container-name namespace namespace-name service service-name time duration
```

```
request paragon debug logs namespace namespace-name service service-name time duration
```

Description

Generate a log file for a service within a namespace. You can also generate this log for a specified period of time.

Options

container <i>container-name</i>	(Optional) Name of the container.
namespace <i>namespace-name</i>	(Mandatory) Name of the namespace. See "request paragon debug logs namespace" on page 1561 for more information.
service <i>service-name</i>	(Mandatory) Name of the service within the namespace. See "request paragon debug logs namespace" on page 1561 for more information.
time <i>duration</i>	Get logs for the specified time period. Format: [0-9]+(s/m/h/d). For example: 5s or 10m

Required Privilege Level

view

Output Fields

When you enter this command, a log file is created for the service and for a specified period of time.

Sample Output

request paragon debug logs namespace healthbot service config-server container config-server time 6s

```
user@node> request paragon debug logs namespace healthbot service config-server container config-server time 6s
Log file for config-server created at: /root/troubleshooting/log/logs-config-server-config-server-2024-04-30_13-42-14.tar.gz
```

request paragon debug logs namespace healthbot service tsdb-shim time 6s

```
user@node> request paragon debug logs namespace healthbot service tsdb-shim time 6s
Log file for tsdb-shim created at: /root/troubleshooting/log/logs-tsdb-shim-2024-04-16_14-49-39.tar.gz
```

request paragon debug logs namespace trust service network

```
user@node> request paragon debug logs namespace trust service network
Log file for network created at: /root/troubleshooting/log/logs-network-2025-06-09_08-59-28.tar.gz
```

Release Information

Command introduced in Juniper Paragon Automation Release 2.0.0.

RELATED DOCUMENTATION

[request paragon debug logs namespace](#) | 1561

request paragon debug logs namespace

SUMMARY

You can generate a log file of various services within a namespace. You can then share this file with the [Juniper Technical Assistance Center \(JTAC\)](#) for further evaluation. Read these topics for more information on the healthbot, foghorn, papi, northstar, and airflow namespaces, and their services.

IN THIS SECTION

- [request paragon debug logs namespace healthbot | 1561](#)
- [request paragon debug logs namespace foghorn | 1563](#)
- [request paragon debug logs namespace papi | 1565](#)
- [request paragon debug logs namespace northstar | 1567](#)
- [request paragon debug logs namespace airflow | 1569](#)

request paragon debug logs namespace healthbot

IN THIS SECTION

- [Syntax | 1562](#)
- [Description | 1562](#)
- [Options | 1562](#)
- [Required Privilege Level | 1562](#)
- [Output Fields | 1563](#)
- [Sample Output | 1563](#)
- [Release Information | 1563](#)

Syntax

```
request paragon debug logs namespace healthbot service service-name
```

```
request paragon debug logs namespace healthbot service service-name time duration
```

Description

Generate a log file for a particular service within the healthbot namespace.

Options

service tsdb-shim	Generate a log file of the tsdb-shim service.
service tand	Generate a log file of the tand service.
service jtimon	Generate a log file of the jtimon service.
service config-server	Generate a log file of the config-server service.
service api-server	Generate a log file of the api-server service.
service analytical-engine	Generate a log file of the analytical-engine service.
service alerta	Generate a log file of the alerta service.
time <i>duration</i>	Get logs for the specified time period. Format: [0-9]+(s/m/h/d). For example: 5s or 10m

Required Privilege Level

view

Output Fields

When you enter this command, a log file is created for the service. The path to the log file is also provided.

Sample Output

request paragon debug logs namespace healthbot service tsdb-shim

```
user@node> request paragon debug logs namespace healthbot service tsdb-shim
```

```
Log file for tsdb-shim created at: /root/troubleshooting/log/logs-tsdb-  
shim-2024-04-16_06-59-58.tar.gz
```

request paragon debug logs namespace healthbot service tsdb-shim time 6s

```
user@node> request paragon debug logs namespace healthbot service tsdb-shim time 6s
```

```
Log file for tsdb-shim created at: /root/troubleshooting/log/logs-tsdb-  
shim-2024-04-16_14-49-39.tar.gz
```

Release Information

Command introduced in Juniper Paragon Automation Release 2.0.0.

request paragon debug logs namespace foghorn

IN THIS SECTION

- [Syntax | 1564](#)
- [Description | 1564](#)
- [Options | 1564](#)
- [Required Privilege Level | 1564](#)
- [Output Fields | 1564](#)
- [Sample Output | 1565](#)
- [Release Information | 1565](#)

Syntax

```
request paragon debug logs namespace foghorn service service-name
```

```
request paragon debug logs namespace foghorn service service-name time duration
```

Description

Create a log file for a particular service within the foghorn namespace.

Options

service order-management	Generate a log file of the order-management service.
service service placement	Generate a log file of the placement service.
service cmgd	Generate a log file of the cmgd service.
service workflow-manager	Generate a log file of the workflow-manager service.
time <i>duration</i>	Get logs for the specified time period. Format: [0-9]+(s/m/h/d). For example: 5s or 10m

Required Privilege Level

view

Output Fields

When you enter this command, a log file is generated for the service. The path to the log file is also provided.

Sample Output

request paragon debug logs namespace foghorn service cmgd

```
user@node> request paragon debug logs namespace foghorn service cmgd
Log file for cmgd created at: /root/troubleshooting/log/logs-cmgd-2024-04-16_07-04-36.tar.gz
```

request paragon debug logs namespace foghorn service cmgd time 6s

```
user@node> request paragon debug logs namespace foghorn service cmgd time 6s
Log file for cmgd created at: /paragon/troubleshooting/log/logs-cmgd-2024-04-30_14-06-41.tar.gz
```

Release Information

Command introduced in Juniper Paragon Automation Release 2.0.0.

request paragon debug logs namespace papi

IN THIS SECTION

- [Syntax | 1566](#)
- [Description | 1566](#)
- [Options | 1566](#)
- [Required Privilege Level | 1566](#)
- [Output Fields | 1566](#)
- [Sample Output | 1567](#)
- [Release Information | 1567](#)

Syntax

```
request paragon debug logs namespace papi service service-name
```

```
request paragon debug logs namespace papi service service-name time duration
```

Description

Generate a log file for a particular service within the papi namespace.

Options

service oc-term	Generate a log file of the oc-term service.
service papi	Generate a log file of the papi service.
service papi-ws	Generate a log file of the papi-ws service.
time <i>duration</i>	Get logs for the specified time period. Format: [0-9]+(s/m/h/d). For example: 5s or 10m

Required Privilege Level

view

Output Fields

When you enter this command, a log file is generated for the service. The path to the log file is also provided.

Sample Output

request paragon debug logs namespace papi service *papi-ws*

```
user@node> request paragon debug logs namespace papi service papi-ws
Log file for papi-ws created at: /root/troubleshooting/log/logs-papi-
ws-2024-04-16_07-14-35.tar.gz
```

request paragon debug logs namespace papi service papi-ws time 2s

```
user@node> request paragon debug logs namespace papi service papi-ws time 2s
Log file for papi-ws created at: /paragon/troubleshooting/log/logs-papi-
ws-2024-04-30_14-05-45.tar.gz
```

Release Information

Command introduced in Juniper Paragon Automation Release 2.0.0.

request paragon debug logs namespace northstar

IN THIS SECTION

- [Syntax | 1568](#)
- [Description | 1568](#)
- [Options | 1568](#)
- [Required Privilege Level | 1568](#)
- [Output Fields | 1568](#)
- [Sample Output | 1569](#)
- [Release Information | 1569](#)

Syntax

```
request paragon debug logs namespace northstar service service-name
```

```
request paragon debug logs namespace northstar service service-name time duration
```

Description

Generate a log file for a service within the northstar namespace.

Options

service toposerver	Generate a log file of the toposerver service.
service configmonitor	Generate a log file of the configmonitor service.
service web	Generate a log file of the web service.
service api-aggregator	Generate a log file of the api-aggregator service.
time <i>duration</i>	Get logs for the specified time period. Format: [0-9]+(s/m/h/d). For example: 5s or 10m

Required Privilege Level

view

Output Fields

When you enter this command, a log file is generated for the service. The path to the log file is also provided.

Sample Output

request paragon debug logs namespace northstar service toposerver

```
user@node> request paragon debug logs namespace northstar service toposerver
```

Log file for toposerver created at: /root/troubleshooting/log/logs-toposerver-2024-04-30_07-33-59.tar.gz

request paragon debug logs namespace northstar service toposerver time 3s

```
user@node> request paragon debug logs namespace northstar service toposerver time 3s
```

Log file for toposerver created at: /paragon/troubleshooting/log/logs-toposerver-2024-04-30_14-10-03.tar.gz

Release Information

Command introduced in Juniper Paragon Automation Release 2.0.0.

request paragon debug logs namespace airflow

IN THIS SECTION

- [Syntax | 1570](#)
- [Description | 1570](#)
- [Options | 1570](#)
- [Required Privilege Level | 1570](#)
- [Output Fields | 1570](#)
- [Sample Output | 1570](#)
- [Release Information | 1571](#)

Syntax

```
request paragon debug logs namespace airflow service service-name
```

```
request paragon debug logs namespace airflow service service-name time duration
```

Description

Generate a log file for a particular service within the airflow namespace.

Options

service workflow-manager	Generate a log file of the workflow-manager service.
time <i>duration</i>	Get logs for the specified time period. Format: [0-9]+(s/m/h/d). For example: 5s or 10m

Required Privilege Level

view

Output Fields

When you enter this command, a log file is generated for the service. The path to the log file is also provided.

Sample Output

request paragon debug logs namespace airflow service workflow-manager

```
user@node> request paragon debug logs namespace airflow service workflow-manager
Log file for workflow-manager created at: /root/troubleshooting/log/logs-workflow-
manager-2024-04-16_14-14-15.tar.gz
```

request paragon debug logs namespace airflow service workflow-manager time 5s

```
user@node> request paragon debug logs namespace airflow service workflow-manager time 5s
```

Log file for workflow-manager created at: /paragon/troubleshooting/log/logs-workflow-manager-2024-04-30_14-11-54.tar.gz

Release Information

Command introduced in Juniper Paragon Automation Release 2.0.0.

request paragon debug postgres

IN THIS SECTION

- [Syntax | 1571](#)
- [Description | 1571](#)
- [Options | 1572](#)
- [Required Privilege Level | 1572](#)
- [Output Fields | 1572](#)
- [Sample Output | 1572](#)
- [Release Information | 1572](#)

Syntax

```
request paragon debug postgres database database-name username user-name measurement measurement-name output file
```

Description

Generate a text file with Postgres information.

Options

database <i>database-name</i>	(Mandatory) Name of the Postgres database.
username <i>user-name</i>	Username of the Postgres database.
measurement <i>measurement-name</i>	Name of the Postgres measurement.
output <i>file</i>	Redirect output to a file.

Required Privilege Level

view

Output Fields

When you enter this command, a text file is generated for Postgres. The path to the text file is also provided.

Sample Output

**request paragon debug postgres database postgres username healthbot measurement
db_counters output file**

```
user@node> request paragon debug postgres database postgres username healthbot measurement  
db_counters output file  
Writing postgres output at: /root/troubleshooting/postgres/postgres_2024-04-16_13-56-27.txt
```

Release Information

Command introduced in Juniper Paragon Automation Release 2.0.0.

request paragon debug redis

IN THIS SECTION

- [Syntax | 1573](#)
- [Description | 1573](#)
- [Options | 1573](#)
- [Required Privilege Level | 1574](#)
- [Output Fields | 1574](#)
- [Release Information | 1574](#)

Syntax

```
request paragon debug redis redis-key-pattern "pattern"
```

```
request paragon debug redis redis-key-string "string | list | set | hash"
```

```
request paragon debug redis redis-key-pattern "pattern" output file
```

Description

Generate Redis key value information and export data to an output file.

Options

redis-key-pattern " <i>pattern</i> "	Display value of Redis key in a specific pattern.
redis-key-string *string list set hash*	Display Redis key value information in the string, list, hash, or set key string types.

output <i>file</i>	Redirect output to a file.

Required Privilege Level

view

Output Fields

request paragon debug redis redis-key-string "test"

```
user@node> request paragon debug redis redis-key-string "test"
REDIS KEY: test, REDIS KEY TYPE: string
REDIS VALUE:
testVal1
```

request paragon debug redis redis-key-pattern "test*"

```
user@node> request paragon debug redis redis-key-pattern "test*"
REDIS KEY: b'test-pattern', REDIS KEY TYPE: string
REDIS VALUE:
testVal2
REDIS KEY: b'test', REDIS KEY TYPE: string
REDIS VALUE:
testVal1
```

Release Information

Command introduced in Juniper Paragon Automation Release 2.0.0.

Service Orchestration

IN THIS CHAPTER

- About the Service Orchestration cMGD CLI | 1576
- set foghorn:core org-id | 1583
- set service design default version | 1584
- set dag state failed | 1586
- show service order network consumers | 1587
- show service order network resources | 1589
- show service order status | 1591
- show service order as-json | 1594
- show service order as-yaml | 1597
- show service designs | 1601
- show device dependent configuration | 1604
- show insights configuration | 1607
- show configuration foghorn:customers | 1610
- show service instance lock status | 1611
- show service instance lock token | 1614
- show service catalog projects | 1616
- show service projects | 1618
- show service order history | 1619
- show airflow dags | 1621
- show airflow dag-runs | 1623
- show paragon orgs | 1626
- request service project add | 1627
- request service catalog project add | 1629
- request service order sync | 1630
- request network resources load | 1631
- request service order load | 1633

- [request service order place | 1634](#)
- [request service order modify | 1636](#)
- [request service catalog project | 1637](#)
- [request service catalog design delete | 1639](#)
- [request service order delete | 1640](#)
- [request service order submit | 1641](#)
- [request service order provision | 1642](#)
- [request service design install | 1643](#)
- [request service design uninstall | 1644](#)
- [request create graphdb | 1646](#)
- [request service order set | 1647](#)
- [request service instance modify design-version | 1648](#)
- [request service instance lock | 1649](#)
- [request service instance unlock | 1651](#)

About the Service Orchestration cMGD CLI

IN THIS SECTION

- [Access the Service Orchestration cMGD CLI | 1577](#)
- [Directories in the Service Orchestration cMGD | 1578](#)
- [Service Orchestration cMGD CLI Commands | 1578](#)

Routing Director provides the service orchestration Containerized Management Daemon (cMGD) CLI that allows root users to perform certain tasks for service provisioning. Before accessing the service orchestration cMGD CLI, you must deploy your cluster and install Routing Director on the cluster. After you log in to the Paragon Shell CLI as the root user, you must exit the Paragon Shell CLI and access the service orchestration cMGD CLI.



NOTE: Only a root user can access the service orchestration cMGD CLI and execute commands for service provisioning tasks.

Access the Service Orchestration cMGD CLI

To access the service orchestration cMGD CLI and execute service orchestration CLI commands, perform the following steps:

1. Log in to the Paragon Shell CLI as the root user.
2. Type `exit` to exit Paragon Shell to the Linux root shell.
3. Start a cMGD pod using the following command:

```
root@node:~# kubectl -n foghorn scale --replicas=1 deployments/cmgd
deployment.apps/cmgd scaled
```

4. Verify that the pod is running. For example:

```
root@node:~# kubectl -n foghorn get pods -l app=cmgd
NAME                                READY   STATUS    RESTARTS   AGE
cmgd-5bcb9568c6-hfkdj             1/1     Running   0           44s
```

5. Launch the service orchestration cMGD CLI:

```
root@node:~# kubectl -n foghorn exec -it -c cmgd deployment/cmgd -- cli
root@cmgd-5bcb9568c6-hfkdj>
```

You are now logged in to the service orchestration cMGD CLI.

6. Use the service orchestration CLI commands to execute service provisioning tasks. See [Table 340 on page 1371](#) to view the commands you can use to navigate in the service orchestration cMGD CLI.
7. When you have executed the required service orchestration CLI commands, type `exit` to exit service orchestration cMGD CLI.
8. Terminate the cMGD pod from the Linux root shell.

```
root@node:~# kubectl -n foghorn scale --replicas=0 deployments/cmgd
deployment.apps/cmgd scaled
```

Directories in the Service Orchestration cMGD

The service orchestration cMGD has default directories that contain the service design YANG files, log files, and so on. [Table 355 on page 1578](#) lists some of the important directories in the service orchestration cMGD.

Table 355: Directories in the Service Orchestration cMGD

Directory	Description
/data-models/projects	Projects contain one or more related service designs based on which service orchestration can provision a service. Routing Director provides the following projects and related YANG models: • network-resource.tgz. • onboard.tgz • L3VPN.tgz • L2VPN.tgz
/foghorn/network/projects	Contains the projects that you upload to the service orchestration cMGD environment by using the request service project add command.
/var/tmp/add_project.sh.log	Contains the log file generated when you execute the request service project add command to add a project to the cMGD environment. The add_project.sh.log file is useful for debugging errors that occur while uploading a project.
/var/tmp/order_sync_conf.json	Contains the configurations (in JSON format) that you synchronize from the order manager to the service orchestration cMGD environment by using the request service order sync command. The order_sync_conf.json file is useful for troubleshooting errors that occur while synchronizing.

Service Orchestration cMGD CLI Commands

Root users can use the commands listed in [Table 356 on page 1579](#) to execute service provisioning tasks in Routing Director:

Table 356: Service Orchestration cMGD CLI Commands and Descriptions

Command	Description
<code>set foghorn:core org-id</code>	Set the organization ID in the service orchestration cMGD environment.
<code>set service design default version</code>	Set the default version for a service design.
<code>set dag state failed</code>	Set a workflow run to failed state.
<code>show service order network consumers</code>	Displays information about dependent service instances for a resource instance.
<code>show service order network resources</code>	Displays all network resources for the specified service instance.
<code>show service order status</code>	View the status of all service orders generated for an organization.
<code>show service order as-json</code>	View all or a specific service order in the JSON format.
<code>show service order as-yaml</code>	View all or a specific service order in the YAML format.
<code>show service designs</code>	View the service design catalog installed for an organization.
<code>show device dependent configuration</code>	View the device-centric configurations for provisioning services, in the XML format.
<code>show insights configuration</code>	View the configurations related to Paragon Insights for monitoring the services that Routing Director intends to provision on a device.

Table 356: Service Orchestration cMGD CLI Commands and Descriptions (Continued)

Command	Description
<code>show configuration foghorn:customers</code>	View configurations for all the services provisioned for a customer.
<code>show service instance lock status</code>	See the lock status of a service instance. A service instance is locked when the instance is being uploaded, modified, or if the network administrator has locked it.
<code>show service instance lock token</code>	See the lock token for a service instance. NOTE: Only superusers must execute this command to acquire the secure lock token stored in the Routing Director database.
<code>show service catalog projects</code>	Displays the projects available in the service catalog for all organizations.
<code>show service projects</code>	Shows the projects that are added to the service orchestration cMGD.
<code>show service order history</code>	Shows the history of all service orders.
<code>show airflow dags</code>	View all workflows generated for an organization.
<code>show airflow dag-runs</code>	View the workflow runs generated for a workflow.
<code>show paragon orgs</code>	View all the organizations created in Juniper Routing Director.
<code>request service project add</code>	Add new service designs and the related YANG models to Routing Director.

Table 356: Service Orchestration cMGD CLI Commands and Descriptions *(Continued)*

Command	Description
<code>request service catalog project add</code>	Add new projects to the service catalog for all organizations.
<code>request service order sync</code>	Synchronize a service order from the order manager and store it in the service orchestration cMGD.
<code>request network resources load</code>	Add network resource pools to the Routing Director database.
<code>request service order load</code>	Upload a service order to the service orchestration cMGD environment.
<code>request service order place</code>	Select placement options and create placement configurations for a service.
<code>request service order modify</code>	Execute the modify workflow for a service.
<code>request service catalog project</code>	Download the projects available in the service catalog for all organizations.
<code>request service catalog design delete</code>	Delete an unused version of a service design from the service catalog.
<code>request service order delete</code>	Execute the delete workflow for a service.
<code>request service order submit</code>	Activate the provisioning workflow for a service order.
<code>request service order provision</code>	Create and execute the workflow for a service order.

Table 356: Service Orchestration cMGD CLI Commands and Descriptions *(Continued)*

Command	Description
<code>request service design uninstall</code>	Uninstall a service design version from the Routing Director database.
<code>request service design install</code>	Install a service design version to the database.
<code>request create graphdb</code>	Create a database for an organization in the service orchestration cMGD.
<code>request service order set</code>	Sets the delete operation for a service order in the service orchestration cMGD. This command does not execute the delete workflow for the order.
<code>request service instance modify design-version</code>	Modify the service design version of a service instance.
<code>request service instance lock</code>	Set a lock for a service instance.
<code>request service instance unlock</code>	Unlock a locked service instance.

SEE ALSO

| [Paragon Shell Overview](#) | **1368**

set foghorn:core org-id

IN THIS SECTION

- [Syntax | 1583](#)
- [Description | 1583](#)
- [Options | 1583](#)
- [Release Information | 1584](#)

Syntax

```
set foghorn:core org-id organization-id
```

Description

Set organization ID in service orchestration cMGD to set the context of all tasks that are executed.

Follow these steps to set organization ID using the command:

1. Log in to the service orchestration cMGD CLI. See ["Access the Service Orchestration cMGD CLI" on page 1577](#).
2. Execute `configure` to enter the configuration mode.
3. Execute the `set foghorn:core org-id organization-id` command.
4. Execute `commit` and `quit` to commit the update and exit the configuration mode.

Options

<i>organization-id</i>	The organization ID is an auto-generated alphanumeric ID assigned to an organization. Get the ID of the organization from the Routing Director GUI. To get the organization ID: 1. Click Settings Menu > System Settings on the banner. 2. Click Copy next to the Organization ID field.

Release Information

Command introduced in Juniper Paragon Automation Release 2.0.0.

RELATED DOCUMENTATION

request create graphdb 1646
request service design install 1643
request service design uninstall 1644
set service design default version 1584

set service design default version

IN THIS SECTION

- [Syntax | 1585](#)
- [Description | 1585](#)
- [Options | 1585](#)
- [Release Information | 1585](#)

Syntax

```
set service design default version design-idversion-number
```

Description

Sets the default version of a service design.

Routing Director supports up to three concurrent versions of a service design. If more than one version is available for a service design, you can set one of the versions as the default version. Any service instance you create, modify, or delete is associated with the default version of the service design.

To change the default version of a service design:

**NOTE:** You must set the organization ID before using this command. See "[set foghorn:core org-id](#)" on page 1583.

1. Log in to the service orchestration cMGD CLI. See "[Access the Service Orchestration cMGD CLI](#)" on page 1577.
2. Execute the `show service designs` command to view the service design catalog.
3. Execute the `set service catalog design default version design-id` command to view the available versions of the specified service design that you can set as the default version.
4. Execute the `set service design default design-idversion-number` command to set the default version of the specified service design.

Options

<i>design-id</i>	Name of the service design for which you want to set the default version. For example, vpn.
<i>version-number</i>	Version of the service design that you want to set as the default version. For example, 1.3.0

Release Information

Command introduced in Juniper Paragon Automation Release 2.0.0.

RELATED DOCUMENTATION

set foghorn:core org-id 1583
show service designs 1601

set dag state failed

IN THIS SECTION

- Syntax | 1586
- Description | 1586
- Options | 1586
- Sample Output | 1587
- Release Information | 1587

Syntax

```
set dag state failed dag-id dag-id run-id run-id
```

Description

Set a workflow run to failed state.

You can set a workflow run to failed state if errors in workflow tasks prevent the workflow from being successfully executed without displaying the state as failed. After setting it to failed state, you can reprovision the service order to reactivate the provisioning workflow.

Options

dag-id <i>dag-id</i>	Unique ID generated for a workflow.
run-id <i>run-id</i>	Unique ID assigned to a workflow run.

Sample Output

set dag state failed

```
root@cmgd-66b8c769c4-fbk74> set dag state failed dag-id onboard-infrastructure-0_10_26 run-id
manual__2024-08-23T07:26:20+00:00
DAG ID: onboard-infrastructure-0_10_26
Run ID: manual__2024-08-23T07:26:20+00:00
DAG state changed to: failed
```

Release Information

Command introduced in Juniper Paragon Automation Release 2.2.0.

show service order network consumers

IN THIS SECTION

- [Syntax | 1587](#)
- [Description | 1587](#)
- [Options | 1588](#)
- [Sample Output | 1588](#)
- [Release Information | 1589](#)

Syntax

```
show service order network consumers resource-instance
```

Description

Display information about dependent service instances for a resource instance.

Options

<i>resource-instance</i>	Display the service instances which use the network resources and are dependent on the specified resource-instance.

Sample Output

show service order network consumers

```
root@6e0385b1eca> show service order network consumers ?
Possible completions:
  <resource-instance>
  internal_l2-addr    Show information for service order 'internal_l2-addr'
  internal_l3-addr    Show information for service order 'internal_l3-addr'
  internal_network    Show information for service order 'internal_network'
  internal_routing    Show information for service order 'internal_routing'
  internal_topo       Show information for service order 'internal_topo'
  internal_vpn-resources Show information for service order 'internal_vpn-resources'
  test-1_l3vpn-1      Show information for service order 'test-1_l3vpn-1'

##### The following output shows that internal_network has a resource dependency on
internal_l3-addr #####
root@6e0385b1eca> show service order network consumers internal_l3-addr
ServiceInstance      Type      Timestamp      Status
internal_network      modify    2024-12-06T17:10:09Z  placed

##### The following output shows that test-1_l3vpn-1 has a resource dependency on
internal_vpn-resources #####
root@6e0385b1eca> show service order network consumers internal_vpn-resources
ServiceInstance      Type      Timestamp      Status
test-1_l3vpn-1       create    2024-12-06T17:10:27Z  applied device
                                                                config to
                                                                [LAX2(LAX2),
                                                                LAX1(LAX1)]

##### The following output is empty as no service instances have resource dependency on
test-1_l3vpn-1 #####
root@6e0385b1eca> show service order network consumers test-1_l3vpn-1
```

```
root@6e0385b1eca>
```

Release Information

Command introduced in Juniper Paragon Automation Release 2.3.0.

RELATED DOCUMENTATION

| [show service order network resources](#) | [1589](#)

show service order network resources

IN THIS SECTION

- [Syntax](#) | [1589](#)
- [Description](#) | [1589](#)
- [Options](#) | [1589](#)
- [Sample Output](#) | [1590](#)
- [Release Information](#) | [1591](#)

Syntax

```
show service order network resources service-instance
```

Description

Display all network resources for the specified service instance.

Options

<i>service-instance</i>	Display the network resources used by the specified service instance in the JSON format.

Sample Output

show service order network resources

```
root@6e0385b1eca> show service order network resources test-1_l3vpn-1
{
  "bandwidth": [
    {
      "poolpath": [
        {
          "name": "LAX",
          "type": "pop"
        },
        {
          "name": "LAX1",
          "type": "pe"
        }
      ],
      <output snipped>
      "poolpath": [
        {
          "name": "ethernet",
          "type": "ifType"
        },
        {
          "name": "40000",
          "type": "ifSpeed"
        },
        {
          "name": "ce4",
          "type": "ce"
        },
        {
          "name": "LAX",
```

```

        "type": "pop"
      },
      {
        "name": "LAX2",
        "type": "pe"
      },
      {
        "name": "ge-0/0/2",
        "type": "port"
      }
    ],
  ]
}

```

Release Information

Command introduced in Juniper Paragon Automation Release 2.3.0.

RELATED DOCUMENTATION

| [show service order network consumers](#) | 1587

show service order status

IN THIS SECTION

- [Syntax](#) | 1592
- [Description](#) | 1592
- [Options](#) | 1592
- [Output fields](#) | 1592
- [Sample Output](#) | 1593
- [Release Information](#) | 1594

Syntax

```
show service order status

show service order status service-order-name

show service order status | display XML
```

Description

See the execution status of service orders generated for service instances.

Options

none	Displays the status of all service orders in an organization.
<i>service-order-name</i>	Name of the service instance for which you want to see the execution status. This option allows you to view the status of a specific service order for a service instance.
display XML	Displays the status of all service orders in an organization in the XML format.

Output fields

[Table 357 on page 1593](#) lists the output fields for the `show service order status` command.

Table 357: show service order status Output Fields

Field Name	Description
Service Instance	Name of the service instance and customer for whom the service is provisioned. The name is displayed in the <i>customer-name_service-instance name</i> format.
Type	Type of operation executed on the service instance: • create • modify • delete
Timestamp	Date and time the service order was generated.
Status	Status of the service order initiated for the service instance. For example, for a network resource service order that is executed, the status displays as network resources updated.

Sample Output

show service order status

ServiceInstance	Type	Timestamp	Status
network-operator_l3-stuff	create	2024-04-16T19:33:12Z	network resources updated
network-operator_rds-and-rts	create	2024-04-16T20:40:56Z	network resources updated
network-operator_topology	modify	2024-04-16T21:39:53Z	network resources updated
att_l3vpn	upgrade	2024-04-18T22:36:30Z	applied device config to [00000000-0000-0000-1000-d4996c2d5200,

```

d4996c2d4e00]                                00000000-0000-0000-1000-
att_l3vpn1                                modify    2024-04-17T13:04:06Z    applied device
                                                config to
d4996c2d5200,                                [00000000-0000-0000-1000-
                                                00000000-0000-0000-1000-
d4996c2d4e00]

```

Release Information

Command introduced in Juniper Paragon Automation Release 2.0.0.

show service order as-json

IN THIS SECTION

- [Syntax | 1594](#)
- [Description | 1595](#)
- [Options | 1595](#)
- [Sample Output | 1596](#)
- [Release Information | 1597](#)

Syntax

```
show service order as-json
```

```
show service order as-json service-order-name
```

```
show service order as-json service-order-name | save directory-path/file-name.json
```

Description

Displays details of a specific order in the JSON format. If you do not specify a service order name for viewing details, details of all the service orders generated in an organization is displayed as an array. You can save the service order details in the JSON format in the service orchestration cMGD pod file system by providing the directory and filename to save.

You can use the downloaded service order JSON files as sample templates, edit the values on the files, and upload the files to create new service orders by using the Routing Director GUI, service orchestration cMGD CLI, or REST API.

Options

<code>none</code>	View all service orders generated for an organization in the JSON format, as an array.
<code><i>service-order-name</i></code>	Specify the name of the service order that you want to view in the JSON format. This option allows you to view a specific order in the JSON format.
<code>save</code>	This option allows you to save service orders in the JSON format to the service orchestration cMGD pod file system.
<code><i>directory-path</i></code>	The location in the service orchestration cMGD pod file system where you want to save the JSON file. For example, <code>/var/tmp</code> .
<code><i>file-name</i></code>	File name you assign to the JSON file when you are saving the file in the service orchestration cMGD pod file system.

Sample Output

show service order as-json *service-order-name*

```
[
  {
    "customer_id": "ebd5a8bf-e2bf-11ee-a392-f6e61c3d4879",
    "design_id": "routing",
    "design_version": "0.1.0",
    "fh_config": {
      "edited-by": "test@gmail.com",
      "status": "network resources updated"
    },
    "instance_id": "l3vpn-routing-iid",
    "operation": "create",
    "order_id": "cc198caf-e2c1-11ee-a392-f6e61c3d4879",
    "org_id": "3c97990f-8446-4fa5-9b88-edd3ddad6cb0",
    "routing": {
      "autonomous_system": [
        {
          "count": 1024,
          "name": 65001
        }
      ],
      "route_reflector": {
        "clusters": [
          {
            "cluster": "1.1.1.1"
          },
          {
            "cluster": "2.2.2.2"
          },
          {
            "cluster": "3.3.3.3"
          }
        ]
      },
      "spring": {
        "sids": {
          "size": 1000
        }
      }
    }
  }
]
```

```

    },
    "upload_time": "2024-03-15T11:47:38Z",
    "version": "1.0.0",
    "workflow_run_id": "manual__2024-03-15T11:47:35.218599+00:00"
  }
]

```

Release Information

Command introduced in Juniper Paragon Automation Release 2.0.0.

show service order as-yaml

IN THIS SECTION

- [Syntax | 1597](#)
- [Description | 1598](#)
- [Options | 1598](#)
- [Sample Output | 1598](#)
- [Release Information | 1601](#)

Syntax

```
show service order as-yaml
```

```
show service order as-yaml service-order-name
```

```
show service order as-yaml service-order-name | save directory-path/file-name.yaml
```

Description

Displays details of a service order in the YAML format. If you do not specify a service order name for viewing details, details of all the service orders generated in an organization is displayed as an array. You can save the service order details in the YAML format in the service orchestration cMGD pod file system by providing the directory and filename to save.

Options

none	View all service orders generated for an organization in the YAML format, as an array.
<i>service-order-name</i>	The service order that you want to view in the YAML format. This option allows you to view a specific order in the YAML format.
save	This option allows you to save service orders in the YAML format to the service orchestration cMGD pod file system.
<i>directory-path</i>	The location in the service orchestration cMGD pod file system where you want to save the YAML file. For example, /var/tmp .
<i>file-name</i>	File name you assign to the YAML file when you are saving the file in the service orchestration cMGD pod file system.

Sample Output

show service order as-yaml

```
- "customer_id": "ebd5a8bf-e2bf-11ee-a392-f6e61c3d4879"
  "design_id": "routing"
  "design_version": "0.1.0"
  "fh_config":
    "edited-by": "test@gmail.com"
```

```

    "status": "network resources updated"
  "instance_id": "l3vpn-routing-iid"
  "operation": "create"
  "order_id": "cc198caf-e2c1-11ee-a392-f6e61c3d4879"
  "org_id": "3c97990f-8446-4fa5-9b88-edd3ddad6cb0"
  "routing":
    "autonomous_system":
      - "count": !!int "1024"
        "name": !!int "65001"
    "route_reflector":
      "clusters":
        - "cluster": "1.1.1.1"
        - "cluster": "2.2.2.2"
        - "cluster": "3.3.3.3"
      "spring":
        "sids":
          "size": !!int "1000"
  "upload_time": "2024-03-15T11:47:38Z"
  "version": "1.0.0"
  "workflow_run_id": "manual__2024-03-15T11:47:35.218599+00:00"
- "customer_id": "ebd5a8bf-e2bf-11ee-a392-f6e61c3d4879"
  "design_id": "vpn"
  "design_version": "0.3.0"
  "fh_config":
    "edited-by": "test@gmail.com"
    "status": "network resources updated"
  "instance_id": "l3vpn-vpn-iid"
  "operation": "create"
  "order_id": "d66e57c1-e2c1-11ee-a392-f6e61c3d4879"
  "org_id": "3c97990f-8446-4fa5-9b88-edd3ddad6cb0"
  "upload_time": "2024-03-15T11:47:53Z"
  "version": "1.0.0"
  "vpn":
    "route_distinguisher":
      - "count": !!int "1024"
        "name": !!int "1234"
      - "count": !!int "1024"
        "name": !!int "1235"
    "route_target":
      - "count": !!int "1024"
        "name": !!int "1234"
      - "count": !!int "1024"
        "name": !!int "1235"

```



```

    "workflow_run_id": "manual__2024-03-15T11:47:49.885346+00:00"
  - "customer_id": "ebd5a8bf-e2bf-11ee-a392-f6e61c3d4879"
    "design_id": "topo"
    "design_version": "0.5.4"
    "fh_config":
      "edited-by": "test@gmail.com"
      "status": "network resources updated"
    "instance_id": "l3vpn-topology-iid"
    "operation": "create"
    "order_id": "6f80e20e-e2c2-11ee-a392-f6e61c3d4879"
    "org_id": "3c97990f-8446-4fa5-9b88-edd3ddad6cb0"
    "topo":
      "pop":
        - "name": "f883d678-c42f-4ddc-9476-ba30a1a2ef45"
          "pe":
            - "access":
                - "ce": "ce1"
                  "name": "xe-1/0/3"
                  "speed": !!int "10000"
                  "type": "ethernet"
                  "bandwidth": !!int "40000000"
                  "mac_addrs": !!int "1000000"
                  "name": "00000000-0000-0000-1000-d4996c636668"
                  "routes": !!int "100000"
                "postal_code": "10321"
                "postal_code_matches":
                  - "name": "SVL"
                    "regex": "10..."
            - "name": "f2d77dde-682b-4e16-94a2-97cff5e30fc7"
              "pe":
                - "access":
                    - "ce": "ce2"
                      "name": "xe-1/0/3"
                      "speed": !!int "10000"
                      "type": "ethernet"
                      "bandwidth": !!int "40000000"
                      "mac_addrs": !!int "1000000"
                      "name": "00000000-0000-0000-1000-d4996c4c2ac0"
                      "routes": !!int "100000"
                    "postal_code": "20321"
                    "postal_code_matches":
                      - "name": "BNG"
                        "regex": "20..."

```

```
"upload_time": "2024-03-15T11:52:12Z"
"version": "1.0.0"
"workflow_run_id": "manual__2024-03-15T11:52:07.069486+00:00"
```

Release Information

Command introduced in Juniper Paragon Automation Release 2.0.0.

show service designs

IN THIS SECTION

- [Syntax | 1601](#)
- [Description | 1601](#)
- [Options | 1602](#)
- [Output Fields | 1602](#)
- [Sample Output | 1603](#)
- [Release Information | 1604](#)

Syntax

```
show service designs
```

Description

Displays the catalog of service designs that Routing Director provides. A service design is a template for defining a service. For example, a service design for the VPN network resource includes placeholders for defining route distinguishers and route targets.

The service catalog includes service designs for onboarding devices, creating network resources pools (such as, IP addresses, AS numbers, route distinguishers, route targets, and so on) in Routing Director, and provisioning network services.

Options

This command has no options.

Output Fields

[Table 358 on page 1602](#) lists the output fields for the `show service catalog design` command.

Table 358: show service designs Output Fields

Field Name	Field Description
Design	Name of the service design.
Type	The type of service that the service design provides guidelines for: • Onboarding—Service type for adding network infrastructure; for example, onboarding devices to Routing Director. • Network resource—Service type for adding routing, topology, L3 address, and VPN resource pools to Routing Director. • L3VPN—Service type for provisioning L3VPN service in a network. • L2VPN—Service type for provisioning L2VPN and L2 Circuit services in a network.
Default Version	Service design version that is set as the default version. Routing Director supports up to three concurrent versions of a service design. If more than one installed version is available for a service design, you can set one of the versions as the default version. Any service instance you create, modify, or delete is associated with the default service design version.
Latest Version	Latest version of the service design available in Routing Director.

Table 358: show service designs Output Fields (Continued)

Field Name	Field Description
Version	The installed version of the service design.
Time	The time when the service design version was installed.
User	The user who installed the service design version.

 **NOTE:** If there is more than one installed version of a service design, the Version, Time, and User fields are displayed for each installed version of the service design.

Sample Output

show service designs

Design	Type	DefaultVersion	LatestVersion
l3-addr	network-resource	0.2.0	0.2.0
Version: 0.2.0, Time: 2024-03-15T08:25:48.056152526Z, User: test@gmail.com			
routing	network-resource	0.1.0	0.1.0
Version: 0.1.0, Time: 2024-03-15T08:25:48.105301762Z, User: test@gmail.com			
infrastruct onboard		0.9.17	0.9.17
Version: 0.9.17, Time: 2024-03-15T08:25:56.820331885Z, User: test@gmail.com			
l3vpn	L3VPN	0.3.28	0.3.29
Version: 0.3.29, Time: 2024-03-15T08:26:31.051326523Z, User: test@gmail.com			
Version: 0.3.28, Time: 2024-03-19T10:38:41.76285894Z, User: nobody			
topo	network-resource	0.5.4	0.5.4
Version: 0.5.4, Time: 2024-03-15T08:25:48.159106464Z, User: 3test@gmail.com			
vpn	network-resource	0.3.0	0.3.0
Version: 0.3.0, Time: 2024-03-15T08:25:48.208596733Z, User: test@gmail.com			

Release Information

Command introduced in Juniper Paragon Automation Release 2.0.0.

RELATED DOCUMENTATION

[request service design install | 1643](#)

[set service design default version | 1584](#)

[request service design uninstall | 1644](#)

[Service Design Overview | 651](#)

[About the Service Designs Page | 653](#)

show device dependent configuration

IN THIS SECTION

- [Syntax | 1604](#)
- [Description | 1605](#)
- [Options | 1605](#)
- [Sample Output | 1605](#)
- [Release Information | 1607](#)

Syntax

```
show device dependent configuration device-name
```

```
show device dependent configuration device-name | save directory-path/filename.xml
```

Description

Displays the device-centric configurations for provisioning services, in the XML format. This device-centric information is useful for troubleshooting and debugging errors when a service order associated with a service instance fails to execute.

You can save the device configurations file to the service orchestration cMGD pod file system. To troubleshoot errors on devices or nodes, you must copy the configurations file from the service orchestration cMGD pod to the controller node. To copy the configurations file to the controller node, log in to the controller node and execute the following command:

```
kubect1 cp foghorn/cmgd-pod-name:directory-path/filename.xml filename.xml
```

Options

<i>device-name</i>	Name of the device for which you want to view the configurations for provisioning a service.
save	This option allows you to save the device configurations to the service orchestration cMGD pod file system.
<i>directory-path</i>	The location in the service orchestration cMGD pod file system where you want to save the device configurations file. For example, /var/tmp .
<i>filename.xml</i>	File name you assign to the device configurations file (in XML format) when you are saving the file in the service orchestration cMGD pod file system.

Sample Output

show device dependent configuration device-name

```
<?xml version="1.0" ?>
<devices>
  <device>
    <name>00000000-0000-0000-1000-d4996c4c2ac0</name>
    <configuration>
```

```

<groups replace="replace">
  <name>paragon-service-orchestration</name>
  <chassis>
    <network-services>enhanced-ip</network-services>
  </chassis>
  <firewall>
    <policer>
      <name>ebd5a8bf-e2bf-11ee-a392-f6e61c3d4879-VPN-1-BNG-access-input</name>
      <logical-interface-policer/>
      <if-exceeding>
        <bandwidth-limit>1000000</bandwidth-limit>
        <burst-size-limit>75000</burst-size-limit>
      </if-exceeding>
      <then>
        <discard/>
      </then>
    </policer>
    <policer>
      <name>ebd5a8bf-e2bf-11ee-a392-f6e61c3d4879-VPN-1-BNG-access-output</name>
      <logical-interface-policer/>
      <if-exceeding>
        <bandwidth-limit>1000000</bandwidth-limit>
        <burst-size-limit>75000</burst-size-limit>
      </if-exceeding>
      <then>
        <discard/>
      </then>
    </policer>
  </firewall>
  <interfaces>
    <interface>
      <name>xe-1/0/3</name>
      <mtu>1514</mtu>
      <unit>
        <name>0</name>
        <description>Network access to VPN VPN-1 for customer ebd5a8bf-e2bf-11ee-a392-
f6e61c3d4879</description>
        <family>
          <inet>
            <policer>
              <input>ebd5a8bf-e2bf-11ee-a392-f6e61c3d4879-VPN-1-BNG-access-input</input>
              <output>ebd5a8bf-e2bf-11ee-a392-f6e61c3d4879-VPN-1-BNG-access-output</output>
            </policer>
          </inet>
        </family>
      </unit>
    </interface>
  </interfaces>
</groups>

```

```
<address>  
  <name>40.1.1.1/30</name>  
</address>  
</inet>  
</family>  
</unit>  
</interface>  
</interfaces>
```

Release Information

Command introduced in Juniper Paragon Automation Release 2.0.0.

show insights configuration

IN THIS SECTION

- [Syntax | 1607](#)
- [Description | 1607](#)
- [Options | 1608](#)
- [Sample Output | 1608](#)
- [Release Information | 1609](#)

Syntax

```
show insights configuration
```

Description

Displays the configurations applied to Paragon Insights for all the service instances provisioned in an organization.

Options

This command has no options.

Sample Output

show insights configuration

```
{
  "insights_configurations": [
    {
      "ebd5a8bf-e2bf-11ee-a392-f6e61c3d4879_l3vpn-onboard-iid": [
        {
          "insights_node_config": {
            "device-group-payload": {
              "device-group-name": "f883d678-c42f-4ddc-9476-ba30a1a2ef45",
              "devices": [
                "d4996c636668"
              ],
            },
            "open-config": {
              "gnmi": {
                "enable": true,
                "encoding": "protobuf"
              },
              "initial-sync": true
            },
          },
          "org-id": "3c97990f-8446-4fa5-9b88-edd3ddad6cb0",
          "playbooks": [
            "onboard-interface-kpis",
            "onboard-chassis-kpis",
            "infrastructure-system-kpis",
            "infrastructure-fpc-kpis",
            "infrastructure-ospf-kpis",
            "infrastructure-bgp-kpis",
            "infrastructure-ldp-kpis",
            "infrastructure-rib-fib-kpis"
          ],
          "site-id": "f883d678-c42f-4ddc-9476-ba30a1a2ef45"
        },
        {
          "device-payload": {
            "device-id": "d4996c636668",
            "host": "d4996c636668",
```

```

        "uuid": "d4996c636668",
        "variable": [
            {
                "instance-id": "ebd5a8bf-e2bf-11ee-a392-f6e61c3_l3vpn-
onboard-iid",
                "playbook": "onboard-interface-kpis",
                "rule": "interfaces/check-interface-in-out-errors-traffic-
state-flaps",
                "variable-value": [
                    {
                        "name": "errors-threshold-variable",
                        "value": "1"
                    },
                    {
                        "name": "flaps-threshold-variable",
                        "value": "1"
                    },
                    {
                        "name": "interface-name",
                        "value": "^xe-1/0/4$"
                    },
                    {
                        "name": "traffic-high-threshold",
                        "value": "80"
                    },
                    {
                        "name": "traffic-low-threshold",
                        "value": "50"
                    }
                ]
            }
        ],
    },

```

Release Information

Command introduced in Juniper Paragon Automation Release 2.0.0.

show configuration foghorn:customers

IN THIS SECTION

- [Syntax | 1610](#)
- [Description | 1610](#)
- [Options | 1610](#)
- [Sample Output | 1610](#)
- [Release Information | 1611](#)

Syntax

```
show configuration foghorn:customers customer-id
```

Description

Displays all the available service configurations for a customer in an organization.

Options

<i>customer-id</i>	Name of the customer.

Sample Output

```
customer f8d24492-f9b1-11ee-a6db-567884529ab3 {
  services L3VPN-testing {
    infrastructure {
      placement {
        infrastructure-ntw {
          network-nodes {
            network-node epic-240-02 {
              ipv4-loopback {
```

```

        selected 192.168.2.0/24;
        options [ 192.168.1.0/24 192.168.2.0/24 192.168.3.0/24 ];
    }
    interfaces {
        interface xe-1/0/1.0 {
            link PE1-----PE2;
            ipv4-address {
                selected 10.1.1.2/30;
                options [ 10.10.1.0/30 10.20.1.0/30 10.30.1.0/30
10.40.1.0/30 10.50.1.0/30 10.60.1.0/30 10.70.1.0/30 ];
            }
        }
    }
}
...

```

Release Information

Command introduced in Juniper Paragon Automation Release 2.0.0.

show service instance lock status

IN THIS SECTION

- [Syntax | 1612](#)
- [Description | 1612](#)
- [Options | 1612](#)
- [Output Fields | 1612](#)
- [Sample Output | 1613](#)
- [Release Information | 1614](#)

Syntax

```
show service instance lock status
```

```
show service instance lock status service-instance-name
```

```
show service instance lock status expired
```

```
show service instance lock status service-instance-name expired
```

Description

See the lock status of a service instance.

A service instance is locked when the instance is being uploaded, modified, or if the network administrator has manually locked it by using the [request service instance lock](#) command in the service orchestration cMGD.

Options

none	See the lock status of all service instances.
<i>service-instance-name</i>	Name of the service instance for which you want to see the lock status.
expired	See the last lock status of all service instances.
<i>service-instance-name</i> expired	See the last lock status of the service instance you specify.

Output Fields

[Table 359 on page 1613](#) lists the output fields for the `show service instance lock status` command.

Table 359: Output Fields for show service instance lock status Command

Field	Description
Service Instance	Name of the service instance.
Acquired By	Name of the user who has acquired the lock for the service instance.
Start Time	The starting time of the lock.
Release Time	The release time of the lock.
Lock Reason	The reason for locking the service instance. For example, Upload service instance.
Release Reason	The reason for which the lock was released. For example, Upload service instance done. If the service instance is still in locked state, the release reason field shows that the instance is still locked.

Sample Output

show service instance lock status

```

root@cmgd-594bd779bd-qqb56> show service instance lock status
ServiceInstance      AcquiredBy      StartTime      ReleaseTime
internal_vpn         cmgd            2024-06-12T13:07:23.070677060Z
0001-01-01T00:00:00Z      LockReason: cmgd-test ReleaseReason: Not released : Instance
still locked

```

show service instance lock status expired

```

root@cmgd-594bd779bd-qqb56> show service instance lock status
ServiceInstance      AcquiredBy      StartTime      ReleaseTime
test_l3vpn-svc       internal         2024-06-10T06:34:48.516222655Z
2024-06-10T06:34:48.699125206Z
LockReason: Upload service instance

```

```

ReleaseReason: Upload service instance done

internal_l3-addr          internal          2024-06-10T06:36:48.245343930Z
2024-06-10T06:36:48.385791296Z
LockReason: Upload service instance
ReleaseReason: Upload service instance done

internal_routing          internal          2024-06-10T06:36:49.552680841Z
2024-06-10T06:36:49.648960819Z
LockReason: Upload service instance
ReleaseReason: Upload service instance done

internal_topo             test1@test.com   2024-06-11T10:29:38.443930804Z
2024-06-11T10:34:29.679572176Z
LockReason: test
ReleaseReason: timeout

```

Release Information

Command introduced in Juniper Paragon Automation Release 2.1.0.

show service instance lock token

IN THIS SECTION

- [Syntax | 1614](#)
- [Description | 1615](#)
- [Options | 1615](#)
- [Output Field | 1615](#)
- [Release Information | 1615](#)

Syntax

```
show service instance lock token service-instance-name
```

Description

See the lock token for a locked service instance.

**NOTE:** Only superusers must execute this command from the service orchestration cMGD to see the secure lock token stored in the Routing Director database. Save the lock token for future use such as unlocking a locked service instance.

Options

<i>service-instance-name</i>	Name of the service instance for which you want to see the lock token.

Output Field

[Table 360 on page 1615](#) lists the output for the `show service instance lock token` command.

Table 360: show service instance lock token Output Field

Field	Description
Lock token	The secure lock token is stored in the Routing Director database. Save the lock token for future use such as unlocking a service instance by using the <code>request service instance unlock</code> command.

Release Information

Command introduced in Juniper Paragon Automation Release 2.1.0.

show service catalog projects

IN THIS SECTION

- [Syntax | 1616](#)
- [Description | 1616](#)
- [Options | 1616](#)
- [Output Fields | 1616](#)
- [Sample Output | 1617](#)
- [Release Information | 1618](#)

Syntax

```
show service catalog projects
```

Description

Displays the projects available in the service catalog for all organizations. Projects contain one or more related service designs.

Options

This command has no options.

Output Fields

[Table 1 on page 1617](#) lists the output fields for the *show service catalog projects* command.

Table 361: show service catalog projects Output Fields

Field	Description
Project Name	Name of the project available in the service catalog to onboard devices, add network resources, and provision different services. • network resource—Displays the service design details to add Layer 3, Layer 2, routing, topology, and VPN resource pools to Routing Director. • onboard—Displays the service design details to onboard devices to Routing Director. • L2VPN—Displays the service design details to provision EVPN, EVPN-VPWS, and L2 circuit services in a network. • L3VPN—Displays the service design details to provision L3VPN services in a network.
Project ID	Unique ID assigned to different projects available in the service catalog.

Sample Output

show service catalog projects

```

root@cmgd-36b9a47c4f06> show service catalog projects
Project Name          Project ID
network-resource      154033a0-ff6d-4740-b1da-aa4fbde211a4
  DesignID: network-design Version: 0.1.5
  DesignID: l2-addr Version: 0.1.5
  DesignID: l3-addr Version: 0.2.6
  DesignID: routing Version: 0.1.9
  DesignID: topo Version: 0.6.17
  DesignID: vpn-resources Version: 0.3.4

onboard               8a9ec9ee-f644-4c09-b484-bad3c4f9eb97
  DesignID: infrastructure Version: 0.13.28

```

```

L2VPN                                824500e1-94a8-4768-85c6-92eef7ba3260
    DesignID: elan-evpn-csm Version: 0.5.8
    DesignID: eline-l2circuit-nsm Version: 0.5.8

L3VPN                                95b27d67-03c2-4689-8638-2f59233be9e5
    DesignID: l3vpn Version: 0.7.8

```

Release Information

Command introduced in Juniper Paragon Automation Release 2.4.0.

show service projects

IN THIS SECTION

- [Syntax | 1618](#)
- [Description | 1618](#)
- [Options | 1618](#)
- [Release Information | 1619](#)

Syntax

```
show service projects
```

Description

Shows the projects that are added to the service orchestration cMGD by using the [request service project add](#) command.

Options

This command has no options.

Release Information

Command introduced in Juniper Paragon Automation Release 2.1.0.

show service order history

IN THIS SECTION

- [Syntax | 1619](#)
- [Description | 1619](#)
- [Options | 1619](#)
- [Output Fields | 1619](#)
- [Sample Output | 1620](#)
- [Release Information | 1621](#)

Syntax

```
show service order history
```

Description

View the history of all service orders.

Options

This command has no options.

Output Fields

[Table 362 on page 1620](#) lists the output fields for the `show service order history` command.

Table 362: show service order history Output Fields

Field	Description
Service Instance	The name of the service instance for which the service order was executed.
Type	The type of operation for which the service order was executed. Service orders are executed when you perform the following operations for a service instance: • Create • Modify • Delete
Timestamp	The date and time when the service order was executed.
Rollback ID	The rollback ID or version number of a service order. The rollback ID for the current service order version is 0, for the previous service order version is 1, for the order version before that is 2, and so on.
Status	Status of the service order initiated for a service instance. For example, when a service instance is saved after creating or modifying it, the status is displayed as Uploaded .

Sample Output

show service order history

```

root@cmgd-68f8484bff-vt4mt> show service order history
ServiceInstance      Type      Timestamp      RollbackID      Status
test_vpn-resources   create    2024-09-10T14:17:39Z      0      update network
resources failed
cust2_test            modify    2024-09-09T11:48:11Z      0      uploaded
cust2_test            modify    2024-09-09T11:44:20Z      1      network
resources updated

```

cust2_test	create	2024-09-09T11:40:51Z	2
uploaded			

Release Information

Command introduced in Juniper Paragon Automation Release 2.1.0.

show airflow dags

IN THIS SECTION

- [Syntax | 1621](#)
- [Description | 1621](#)
- [Options | 1622](#)
- [Output Fields | 1622](#)
- [Sample Output | 1622](#)
- [Release Information | 1623](#)

Syntax

```
show airflow dags
```

```
show airflow dags | display xml
```

Description

View all the workflows generated to upload network resources and provision services for an organization. Routing Director uses automated workflows or series of tasks to:

- Upload or delete network resources to and from the Routing Director database.
- Provision or delete services in the network.

Options

display xml	View all workflows generated for an organization in the XML format.

Output Fields

Table 363 on page 1622 lists the output fields for the show airflow dags command.

Table 363: show airflow dags Output Fields

Field Name	Field Description
DAG-ID	Unique ID generated for a workflow.
Description	Description about the workflow, for example, create L3VPN service.

Sample Output

show airflow dags

```
root@cmgd-66b8c769c4-fbk74> show airflow dags
DAG-ID                                     Description
airflow_logs_cleanup
create-infrastructure-0_10_26              Service orchestration create infrastructure
workflow
create-l2-addr-0_1_4                      Service orchestration create l2 address network
resour
create-l2vpn-0_3_41                       Service orchestration create L2VPN service
create-l2vpn-nsm-0_2_22                   Service orchestration create L2VPN NSM service
create-l3-addr-0_2_4                      Service orchestration create l3 address network
resour
create-l3vpn-0_5_58                       Service orchestration create L3VPN service
```

show airflow dags | display xml

```
root@cmgd-66b8c769c4-fbk74> show airflow dags | display xml
<rpc-reply xmlns:junos="http://xml.juniper.net/junos/24.3I0/junos">
  <airflow-dags>
    <dags>
      <dag-id>airflow_logs_cleanup</dag-id>
      <description/>
    </dags>
    <dags>
      <dag-id>create-infrastructure-0_10_26</dag-id>
      <description>Service orchestration create infrastructure workflow</description>
    </dags>
  </airflow-dags>
  <cli>
    <banner></banner>
  </cli>
</rpc-reply>
```

Release Information

Command introduced in Juniper Paragon Automation Release 2.2.0.

RELATED DOCUMENTATION

[About the Workflows Page](#) | 880

show airflow dag-runs

IN THIS SECTION

- [Description](#) | 1624
- [Options](#) | 1624
- [show airflow dag-runs Output Fields](#) | 1624
- [Sample Output](#) | 1625

Syntax

```
show airflow dag-runs dag-id dag-id
```

```
show airflow dag-runs dag-id dag-id | display xml
```

Description

View the workflow runs generated for a workflow. A workflow run is an instance of executing the series of tasks in a workflow.

Options

dag-id <i>dag-id</i>	Unique ID of the workflow for which the workflow run is executed.
display xml	View all runs for a workflow in the XML format.

show airflow dag-runs Output Fields

[Table 364 on page 1624](#) lists the output fields for the `show airflow dag-runs` command.

Table 364: show airflow dag-runs Output Fields

Field Name	Field Description
DAG-Run-ID	Unique ID assigned to a workflow run.
State	Execution state of a workflow run. For example, when a workflow run is successfully executed, the state is displayed as Success . When a workflow run fails to execute, the state is displayed as Failed .

Sample Output

show airflow dag-runs

```
root@cmgd-66b8c769c4-fbk74> show airflow dag-runs dag-id onboard-infrastructure-0_10_26
DAG-Run-ID                                     State
manual__2024-08-22T06:19:49+00:00             failed
manual__2024-08-22T08:57:23+00:00             failed
manual__2024-08-23T07:26:20+00:00             success
manual__2024-08-23T08:58:37+00:00             failed
manual__2024-08-26T11:58:11+00:00             failed
```

show airflow dag-runs dag-id *dag-id* | display xml

```
root@cmgd-66b8c769c4-fbk74> show airflow dag-runs dag-id onboard-infrastructure-0_10_26 |
display xml
<rpc-reply xmlns:junos="http://xml.juniper.net/junos/24.3I0/junos">
  <airflow-dag-runs>
    <dag-runs>
      <run-id>manual__2024-08-22T06:19:49+00:00</run-id>
      <service-order>internal_vpn</service-order>
      <state>failed</state>
    </dag-runs>
  </airflow-dag-runs>
</cli>
  <banner></banner>
</cli>
</rpc-reply>
```

Release Information

Command introduced in Juniper Paragon Automation Release 2.2.0.

RELATED DOCUMENTATION

[About the Workflow Runs Page](#) | 882

show paragon orgs

IN THIS SECTION

- [Syntax | 1626](#)
- [Description | 1626](#)
- [Options | 1626](#)
- [show paragon orgs Output Fields | 1626](#)
- [Release Information | 1627](#)

Syntax

```
show paragon orgs
```

Description

View all the organizations created in Juniper Routing Director.

Options

This command has no options.

show paragon orgs Output Fields

[Table 365 on page 1626](#) lists the output fields for the show paragon orgs command.

Table 365: show paragon orgs Output Fields

Field Name	Field Description
Org-Name	Name of the organization created in Routing Director.
Org-ID	Unique auto-generated alphanumeric ID assigned to an organization.

Release Information

Command introduced in Juniper Paragon Automation Release 2.2.0.

request service project add

IN THIS SECTION

- [Syntax | 1627](#)
- [Description | 1627](#)
- [Options | 1628](#)
- [Release Information | 1628](#)

Syntax

```
request service project add projectName_projectID
```

```
request service project add projectName_projectID force
```

Description

Adds new projects and related YANG models for an organization to the service orchestration cMGD. Projects contain one or more related service designs based on which Routing Director can provision a service. To add a new project to the service catalog, use the [request service catalog project add](#) command.

You can view all the projects available in the service catalog by executing the `show service catalog projects` command. If you don't add *projectName_projectID* option when executing the `request service project add` command, all projects available in the service catalog are added to the service orchestration cMGD, for an organization.

If an error occurs while adding projects to the service orchestration cMGD, see the `/var/tmp/add_project.sh.log` file to troubleshoot the reason for failure.



NOTE: In releases earlier than release 2.4.0, use the request service project add /data-models/projects/*project-name*.tgz (none | force) command to add new projects and related YANG models for an organization to the service orchestration cMGD.

Options

<i>projectName</i>	Name of the project available in the service catalog to onboard devices, add network resources, and provision different services. • network resource—Displays the service design details to add Layer 3, Layer 2, routing, topology, and VPN resource pools to Routing Director. • onboard—Displays the service design details to onboard devices to Routing Director. • L2VPN—Displays the service design details to provision EVPN, EVPN-VPWS, and L2 circuit services in a network. • L3VPN—Displays the service design details to provision L3VPN services in a network.
<i>projectID</i>	Unique ID assigned to different projects available in the service catalog.
force	This option allows you to ignore constraints that lead to errors while adding projects, and add the specified project to the service orchestration cMGD. Errors while adding projects might occur due to constraints such as the service design version is not the same as the latest version available in the service catalog, service designs on which the current service design depends are not uploaded, and so on.

Release Information

Command introduced in Juniper Paragon Automation Release 2.4.0.

request service catalog project add

IN THIS SECTION

- [Syntax | 1629](#)
- [Description | 1629](#)
- [Options | 1629](#)
- [Release Information | 1629](#)

Syntax

```
request service catalog project add project-file
```

```
request service catalog project add /data-models/projects/project-file.tgz
```

Description

Adds new projects to the service catalog for all organizations. Projects contain one or more related service designs based on which service orchestration can provision a service. Execute the request service project add *projectName_projectID* command to install the project for an individual organization.

Options

<i>project-file</i>	Name of the project file you want to add (in TGZ format) to the service catalog.

Release Information

Command introduced in Juniper Paragon Automation Release 2.3.0.

RELATED DOCUMENTATION

| [request service project add](#) | [1627](#)

request service order sync

IN THIS SECTION

- [Syntax](#) | [1630](#)
- [Description](#) | [1630](#)
- [Options](#) | [1631](#)
- [Release Information](#) | [1631](#)

Syntax

```
request service order sync service-order-name
```

```
request service order sync service-order-name rollback rollback-id
```

Description

Synchronize a service order from the order manager and store it in the service orchestration cMGD environment. The service order is stored in the `foghorn:customers` branch of the service orchestration cMGD environment.



NOTE:

- Ensure that you type the `request service order sync` command every time you make configuration changes to a service order as this command synchronizes the service configurations from the order manager to the service orchestration cMGD.

- If an error occurs while synchronizing, refer to the `/var/tmp/order_sync_conf.json` file for troubleshooting.

The rollback option synchronizes and uploads a specified earlier version of a service order into the service orchestration cMGD. To view service order history and rollback IDs for earlier versions, execute the `show service order history` command.

If you want to restore the configurations in the earlier service order version after synchronizing and uploading the order, you can provision the earlier order version by executing the `request service order provision` command.

Options

<i>service-order-name</i>	Name of the service order that you want to synchronize.
rollback	The rollback option synchronizes and uploads a specified earlier version of a service order into the service orchestration cMGD.
<i>rollback-id</i>	Rollback ID of the earlier service order version that you want to upload to the service orchestration cMGD. The rollback ID of the current service order version is 0.

Release Information

Command introduced in Juniper Paragon Automation Release 2.1.0.

request network resources load

IN THIS SECTION

- Syntax | 1632
- Description | 1632
- Options | 1632

Syntax

```
request network resources load service-order-name
```

```
request network resources load service-order-name lock-token lock-token
```

Description

Uploads newly created or modified network resource configuration files for a service instance to the Routing Director database. For more information about uploading network resources by using this command, see ["Add Network Resources for Provisioning Services" on page 670](#).

The lock-token option allows you to unlock a locked service instance by providing the lock token.

Options

<i>service-order-name</i>	Name of the service order for uploading network resources to Routing Director.
lock-token <i>lock-token</i>	The secure lock token that is displayed when you set a lock for a service instance by using the request service instance lock command.

Release Information

Command introduced in Juniper Paragon Automation Release 2.0.0.

request service order load

IN THIS SECTION

- [Syntax | 1633](#)
- [Description | 1633](#)
- [Options | 1633](#)
- [Release Information | 1634](#)

Syntax

```
request service order load service-order-name
```

```
request service order load service-order-name operation operation-name
```

```
request service order load service-order-name merge
```

```
request service order load service-order-name override
```

Description

Uploads a service order to the service orchestration cMGD. This command does not activate the workflow for the respective operation.

You can use the options available with this command to override or merge the configurations specified in the service order that you are uploading to the existing service instance configurations. **Override** is the default action when you upload a service order.

Options

<i>service-order-name</i>	Name of the service order that you are uploading to the service orchestration cMGD.
operation <i>operation-name</i>	Type of operation of the service order: • create • modify The default operation is create.
merge	This option merges the configurations in the service order you are uploading into the existing service instance.
override	This option overrides or replaces the existing service instance configurations with the configurations in the service order you are uploading. Override is the default action when you upload a service order.

Release Information

Command introduced in Juniper Paragon Automation Release 2.1.0.

request service order place

IN THIS SECTION

- [Syntax | 1635](#)
- [Description | 1635](#)
- [Options | 1635](#)
- [Release Information | 1636](#)

Syntax

```
request service order place service-order-name
```

```
request service order place service-order-name lock-token lock-token
```

```
request service order place service-order-name modifier (delete|keep|options|place|release)
```

Description

Select placement options and create placement configurations for a service. See ["Manage Placement Configurations by Using the Service Orchestration cMGD CLI" on page 770](#).

The lock-token option allows you to unlock a locked service instance by providing the lock token.

You can use this command to generate and modify placement configurations for a service instance.

Options

<i>service-order-name</i>	The name of the service order for which placement configurations are to be generated.
lock-token <i>lock-token</i>	The secure lock token that is displayed when you set a lock for a service instance by using the request service instance lock command.

<code>service-order-name</code> modifier (delete keep options place release)	Modify the placements behavior for a service instance. • delete—Delete all network resource allocations for the service instance. • keep—Allocate the network resources and also do not release the deleted network resources. • options—Assign the available placement options, but don't allocate any resources. • place—Allocate the network resources for the service order. • release—Release the deleted network resources that were retained using the keep option.

Release Information

Command introduced in Juniper Paragon Automation Release 2.0.0.

delete option introduced in Juniper Paragon Automation Release 2.3.0.

RELATED DOCUMENTATION

[request service order sync](#) | 1630

[show service order status](#) | 1591

[request service instance lock](#) | 1649

[Manage Placement Configurations for Service Instances](#) | 768

request service order modify

IN THIS SECTION

● [Syntax](#) | 1637

● [Description](#) | 1637

- [Options | 1637](#)
- [Release Information | 1637](#)

Syntax

```
request service order modify service-order-name
```

Description

Activates the modify workflow for a specified service order.

Options

<i>service-order-name</i>	The name of the service order for which you are activating the modify workflow.

Release Information

Command introduced in Juniper Paragon Automation Release 2.0.0.

request service catalog project

IN THIS SECTION

- [Syntax | 1638](#)
- [Description | 1638](#)
- [Options | 1638](#)
- [Sample Output | 1638](#)
- [Release Information | 1638](#)

Syntax

```
request service catalog project project-ID project-id path /file-path
```

Description

Download the projects available in the service catalog for all organizations. You can download the projects to the service orchestration cMGD pod file system.

Options

<code>project-ID <i>project-id</i></code>	Name of the project available in the service catalog that you want to download.
<code>path /<i>file-path</i></code>	Path in the service orchestration cMGD pod file system where the project is to be downloaded.

Sample Output

request service catalog project

```
root@cmgd-83e041f1b9e3> request service catalog project project-ID network-
resource_3728a4bf-05b5-4fbc-9fd0-d38d46e13590 path /tmp
Downloaded catalog project: network-resource_3728a4bf-05b5-4fbc-9fd0-d38d46e13590
Project has been downloaded to: /tmp
```

Release Information

Command introduced in Juniper Paragon Automation Release 2.4.0.

request service catalog design delete

IN THIS SECTION

- [Syntax | 1639](#)
- [Description | 1639](#)
- [Options | 1639](#)
- [Sample Output | 1640](#)
- [Release Information | 1640](#)

Syntax

```
request service catalog design delete designID-version
```

Description

Delete an unused version of a service design from the service catalog. The service design is uninstalled from all organizations and is not available to install again.

**NOTE:** You cannot delete a service design from the service catalog if there are active instances associated with the selected service design.

Options

<i>designID</i>	Name of the service design that you want to delete.
<i>version</i>	Version of the service design that you want to uninstall.

Sample Output

```
root@cmgd-83e041f1b9e3> request service catalog design delete vpn-resources-0.3.4
Deleting service design from catalog
Design ID: vpn-resources
Design Version: 0.3.4
Delete Workflow ID: manual__2025-02-04T05:00:28.095439+00:00
```

Release Information

Command introduced in Juniper Paragon Automation Release 2.4.0.

RELATED DOCUMENTATION

| [request service design uninstall](#) | [1644](#)

request service order delete

IN THIS SECTION

- [Syntax](#) | [1640](#)
- [Description](#) | [1640](#)
- [Options](#) | [1641](#)
- [Release Information](#) | [1641](#)

Syntax

```
request service order delete service-order-name
```

Description

Activates the delete workflow for a specific service order.

Options

<i>service-order-name</i>	The name of the service order for which you want to activate the delete workflow.

Release Information

Command introduced in Juniper Paragon Automation Release 2.0.0.

request service order submit

IN THIS SECTION

- Syntax | 1641
- Description | 1641
- Options | 1641
- Release Information | 1642

Syntax

```
request service order submit service-order-name
```

Description

Activates the add, modify, or delete workflow of a service order that is already present in service orchestration.

Options

<i>service-order-name</i>	The name of the service order for which you are activating the provisioning workflow.

Release Information

Command introduced in Juniper Paragon Automation Release 2.0.0.

request service order provision

IN THIS SECTION

- Syntax | 1642
- Description | 1642
- Options | 1642
- Release Information | 1643

Syntax

```
request service order provision service-order-name
```

Description

Adds a service order to service orchestration and submits the service order. Use this command to add and execute the workflow for a service order by issuing a single command.

Options

<i>service-order-name</i>	The name of the service order that you want to add to service orchestration and provision.

Release Information

Command introduced in Juniper Paragon Automation Release 2.0.0.

request service design install

IN THIS SECTION

- [Syntax | 1643](#)
- [Description | 1643](#)
- [Options | 1644](#)
- [Release Information | 1644](#)

Syntax

```
request service design install design-id
```

Description

Installs a new version of a service design version in an organization.



NOTE: You must set the organization ID before using this command. See "[set foghorn:core org-id](#)" on page 1583.

Routing Director supports up to three concurrent versions of a service design. You can uninstall an unused service design version for an organization and install a newer version to replace the older version.

Perform the following steps to install a service design version by using this command:

1. Log in to the service orchestration cMGD CLI. See "[Access the Service Orchestration cMGD CLI](#)" on page 1577.
2. Execute the `show service designs` command to see the latest version of the service design available in Routing Director.

- 3. Execute the request service design install *design-id* command to install the latest version of the service design.

Options

<i>design-id</i>	Name of the service design for which you want to install the latest available version. For example, l3vpn.

Release Information

Command introduced in Juniper Paragon Automation Release 2.0.0.

RELATED DOCUMENTATION

set foghorn:core org-id 1583
show service designs 1601
request service design uninstall 1644
set service design default version 1584

request service design uninstall

IN THIS SECTION

- [Syntax | 1645](#)
- [Description | 1645](#)
- [Options | 1645](#)
- [Release Information | 1645](#)

Syntax

```
request service design uninstall design-idversion-number
```

Description

Uninstall a specific version of a service design from an organization.

Routing Director supports up to three concurrent versions of a service design. You can uninstall an unused version for an organization and install a newer version to replace the older version.

 **NOTE:**

- You must set the organization ID before using this command. See ["set foghorn:core org-id" on page 1583](#).
- You cannot uninstall a service design version if there are active service instances associated with the design version. You must delete the service instances before uninstalling the service design.

Options

<i>design-id</i>	Name of the service design that you want to uninstall. For example, l3vpn.
<i>version-number</i>	Version of the service design that you want to uninstall. For example, 0.3.1.

Release Information

Command introduced in Juniper Paragon Automation Release 2.0.0.

RELATED DOCUMENTATION

set foghorn:core org-id 1583
request service design install 1643
request service catalog design delete 1639

request create graphdb

IN THIS SECTION

- [Syntax | 1646](#)
- [Description | 1646](#)
- [Options | 1646](#)
- [Release Information | 1646](#)

Syntax

```
request create graphdb
```

Description

Create a database for an organization in the service orchestration cMGD.

Before creating a database for an organization by using this command, you must set the organization ID in the service orchestration cMGD by executing the [set foghorn:core org-id](#) command.

Options

This command has no options.

Release Information

Command introduced in Juniper Paragon Automation Release 2.1.0.

RELATED DOCUMENTATION

| [set foghorn:core org-id | 1583](#)

request service order set

IN THIS SECTION

- Syntax | 1647
- Description | 1647
- Options | 1647
- Release Information | 1648

Syntax

```
request service order set service-order-name operation delete
```

Description

Sets the delete operation for a service order in the service orchestration cMGD.

This command does not execute the delete workflow for the order. After you set the delete operation for the service order, the service order status is **uploaded**. To see the service order status, use the `show service order status service-order-name` command.

To execute the delete workflow for the service order, use the `request service order submit service-order-name` command.

Options

<i>service-order-name</i>	Name of the service order for which you want to set the delete operation.
operation delete	Sets the delete operation for a specified service order in the service orchestration cMGD. The delete workflow is not activated by this option.

Release Information

Command introduced in Juniper Paragon Automation Release 2.1.0.

RELATED DOCUMENTATION

[show service order status | 1591](#)

[request service order submit | 1641](#)

request service instance modify design-version

IN THIS SECTION

- [Syntax | 1648](#)
- [Description | 1648](#)
- [Options | 1649](#)
- [Sample Output | 1649](#)
- [Release Information | 1649](#)

Syntax

```
request service instance modify design-version service-instance-name service-design-version
```

Description

Modify the service design version of a service instance. You can choose to upgrade your default service design version to a later version or downgrade to an earlier version. To change to a service design version, it must be installed. To view a catalog of service designs available, use the `show service designs` command.



NOTE: You cannot downgrade a service design version using the Routing Director GUI.

Options

<i>service-instance-name</i>	The name of the service instance for which you want to modify the service design.
<i>service-design-version</i>	The service design version to which you want to upgrade or downgrade.

Sample Output

request service instance modify design-version

```
root@cmgd-76d447bd6-f9ld> request service instance modify design-version JNPR_l2ckt 0.5.0
Modifying design and publishing service instance: JNPR_l2ckt
Design ID: eline-l2circuit-nsm
Modified Version : 0.5.0
```

Release Information

Command introduced in Juniper Paragon Automation Release 2.3.0.

RELATED DOCUMENTATION

| [request service design install](#) | [1643](#)

request service instance lock

IN THIS SECTION

- [Syntax](#) | [1650](#)
- [Description](#) | [1650](#)
- [Options](#) | [1650](#)

- [Output Field | 1650](#)
- [Sample Output | 1651](#)
- [Release Information | 1651](#)

Syntax

```
request service instance lock service-instance-name lock-reason lock-reason lock-time lock-time
```

Description

Set a lock for a service instance.

Options

<i>service-instance-name</i>	The name of the service instance that you want to lock.
lock-reason <i>lock-reason</i>	The reason for locking the service instance. For example, Upload service instance.
lock-time <i>lock-time</i>	The duration (in seconds) for which you want to lock the service instance. For example, 300. After the specified duration, the service instance is automatically unlocked. As a superuser, you can unlock a service instance before the lock duration is over by using the request service instance unlock command.

Output Field

[Table 366 on page 1651](#) lists the output for the request service instance lock command.

Table 366: request service instance lock Output Field

Field	Description
Lock token	The secure lock token is stored in the Routing Director database. Save the lock token for future use such as unlocking a service instance by using the request service instance unlock command.

Sample Output

request service instance lock

```
root@cmgd-594bd779bd-qqb56> request service instance lock internal_vpn lock-reason cmgd-test
lock-time 300
Locking service instance: internal_vpn
Lock token: cb6a0544-26f3-11ef-b3ef-764115448f0a_vpn:22:_i-wNXz---_
```

Release Information

Command introduced in Juniper Paragon Automation Release 2.1.0.

RELATED DOCUMENTATION

| [request service instance unlock](#) | [1651](#)

request service instance unlock

IN THIS SECTION

- [Syntax](#) | [1652](#)
- [Description](#) | [1652](#)
- [Options](#) | [1652](#)
- [Sample Output](#) | [1652](#)

Syntax

```
request service instance unlock service-instance-name unlock-reason unlock-reason lock-token
lock-token
```

Description

Unlock a locked service instance.

Options

<i>service-instance-name</i>	Name of the service instance that you want to unlock.
unlock-reason <i>unlock-reason</i>	Reason for unlocking the service instance. For example, test done.
lock-token <i>lock-token</i>	The secure lock token that is displayed when you acquire a service instance lock by using the request service instance lock command.

Sample Output

request service instance unlock

```
root@cmgd-594bd779bd-qqb56> request service instance unlock internal_vpn unlock-reason test-done
lock-token cb6a0544-26f3-11ef-b3ef-764115448f0a_vpn:22:_i-wNXz---_
Unlocking service instance: internal_vpn
```

Release Information

Command introduced in Juniper Paragon Automation Release 2.1.0.

RELATED DOCUMENTATION

| [request service instance lock](#) | 1649