

Paragon Automation User Guide

Published
2024-05-21

RELEASE
23.1

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.

Paragon Automation User Guide

23.1

Copyright © 2024 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 | xxi

Introduction

Overview | 2

About the Paragon Automation User Guide | 2

Paragon Automation Overview | 3

Paragon Pathfinder Overview | 6

Paragon Insights Overview | 9

Paragon Planner Overview | 13

Understand Differences between Paragon Pathfinder and Planner | 15

Paragon Automation GUI Overview | 18

Paragon Automation GUI Menu Overview | 27

Access the Paragon Automation GUI | 35

Single Sign-On Overview | 35

Access the Paragon Automation GUI | 36

About the Dashboard Page | 38

Change Your Password | 45

Reset Your Password | 45

Reset User Password Through System Console | 46

Access the Paragon Planner | 48

Access Paragon Planner Desktop Application | 48

Access Paragon Planner Web Application | 49

Configure SMTP, LDAP, and Portal Settings | 50

Configure SMTP Settings | 50

LDAP Authentication Overview | 52

Configure LDAP Settings | 53

Configure Portal Settings | 55

Manage Users | 57

Users Overview | 57

About the Users Page | 58

Add Users | 62

Edit and Delete Users | 64

 | Edit Users | 65

 | Delete Users | 65

Manage Roles | 66

Roles Overview | 66

About the Roles Page | 67

Add Roles | 69

Edit, Clone, and Delete Roles | 71

 | Edit a Role | 72

 | Clone a Role | 72

 | Delete a Role | 73

Manage User Groups | 74

User Groups Overview | 74

About the User Groups Page | 75

Add User Groups | 77

Edit and Delete User Groups | 79

 | Edit User Groups | 79

 | Delete User Groups | 79

Identity Providers | 81

About the Identity Providers Page | 81

Add Identity Providers | 84

Edit and Delete Identity Providers | 87

 | Edit Identity Providers | 87

2

| Delete Identity Providers | 87

Integration With Anuta ATOM

Integration With Anuta ATOM | 90

Integration with Anuta ATOM Overview | 90

Register Paragon Automation with Keycloak | 93

Add Anuta ATOM as an Identity Provider | 97

Add Anuta ATOM as an External EMS Application | 99

Configure Paragon Automation to Route API Calls to Anuta ATOM | 101

3

Workflows

Base Platform | 105

Onboard and Manage Devices | 105

Paragon Pathfinder | 109

Acquire and View the Network Topology | 109

Add LSPs (Tunnels) | 111

Schedule and Monitor a Maintenance Event | 112

Reroute LSPs Automatically | 113

Add and Check Container LSPs | 116

Paragon Planner | 118

Obtaining and Importing Network Files | 118

Plan Network for Optimum Performance | 119

Simulate Failure Scenarios | 120

Paragon Insights | 124

Monitor and Troubleshoot Device and Network Health | 124

Analyze Root Cause of Network, Device, and Service Issues | 125

Bring Your Own Ingest Default Plug-in Workflow | 127

Bring Your Own Ingest Custom Plug-in Workflow | 127

Generate and View Health Reports | 128

Manage Devices and Network

Devices | 130

Devices Overview | 130

Zero-Touch Provisioning Overview | 132

Configure a DHCP Relay for ZTP | 134

Add Devices Overview | 135

About the Devices Page | 138

Supported Devices | 144

Add Devices | 149

Discover Devices | 149

Add New Devices | 153

Upgrade the Device Image | 156

View Device Statistics and Inventory information | 157

View Device Statistics | 157

View Device Inventory | 159

View and Manage Device Configuration | 165

Edit Devices | 168

Delete Devices | 173

Device Groups | 175

About the Device Groups Page | 175

Add a Device Group | 177

Edit a Device Group | 183

Filter a Device Group | 183

Delete a Device Group | 184

Commit or Roll Back Configuration Changes in Paragon Insights | 185

Device Images | 188

Image Upgrade Workflow | 188

About the Images Page | 190

Upload an Image | 192

Stage an Image | 193

Deploy an Image | 194

Delete Images | 196

Network | 197

Assign Names to Admin Group Bits | 197

Modify Pathfinder Settings From the Pathfinder CLI | 198

Access the Pathfinder CLI | 199

Modify Pathfinder Configuration Settings | 203

Modify Pathfinder Settings From the GUI | 206

Disaster Recovery Overview | 252

Network Slicing Overview | 256

Add a Test Agent for Network Slices | 259

Configure LSP Routing in a Network Slice by Using a Path Computation Profile | 261

Network Groups | 264

About the Network Groups Page | 264

Add a Network Group | 266

Edit a Network Group | 269

Topology Filter | 270

About the Topology Filter Page | 270

Add a Topology Filter Rule | 274

Edit a Topology Filter Rule | 275

Delete a Topology Filter Rule | 276

Manage Device Templates and Configuration Templates

Configuration Templates | 278

Configuration Templates Overview | 278

Configuration Templates Workflow	280
About the Configuration Templates Page	281
Add Configuration Templates	284
Preview and Render a Configuration Template	291
Assign Configuration Templates to a Device Template	292
Deploy a Configuration Template to a Device	293
Edit, Clone, and Delete a Configuration Template	295
Edit a Configuration Template	295
Clone a Configuration Template	296
Delete a Configuration Template	296
Device Templates 	298
Device Templates Overview	298
About the Device Templates Page	299
Edit Configuration Templates Assigned to a Device Template	302
Edit, Clone, and Delete Device Templates	303
Edit a Device Template	304
Clone a Device Template	304
Delete Device Templates	305
6	Manage Playbook, Rules, Resources, and Graphs
Playbooks 	307
About Playbooks	307
Add a Predefined Playbook	308
Create a Playbook Using the Paragon Insights GUI	309
Edit a Playbook	310
Clone a Playbook	311
Manage Playbook Instances	312
View Information About Playbook Instances	312
Create and Run a Playbook Instance	314
Manually Pause or Play a Playbook Instance	316

| Create a Schedule to Automatically Play/Pause a Playbook Instance | 317

Rules | 320

Understand Paragon Insights Topics | 320

Rules Overview | 321

About the Rules Page | 340

Add a Predefined Rule | 340

Edit, Clone, Delete, and Download Rules | 341

Configure a Custom Rule in Paragon Automation GUI | 343

 Create a New Rule Using the Paragon Automation GUI | 343

 Rule Filtering | 345

 Sensors | 346

 Fields | 348

 Vectors | 350

 Variables | 352

 Functions | 353

 Triggers | 355

 Rule Properties | 357

 Pre-Action Tasks | 358

 | Post-Action Tasks | 358

Configure Paragon Insights Notification for LSP Gray Failures | 359

Configure Multiple Sensors per Device | 362

Understand Sensor Precedence | 364

Configure Sensor Precedence | 365

Resources | 369

Understand Root Cause Analysis | 369

About the Resources Page | 372

Add Resources for Root Cause Analysis | 375

Configure Dependency Between Resources | 378

Example Configuration: OSPF Resource Dependency | 383

Edit Resources and Dependencies | 394

 | Edit a Resource | 394

 | Edit Resource Dependency | 395

Upload Resources | 395

Download Resources | 396

Clone Resources | 397

Delete User-Generated Resources and Dependencies | 398

 | Delete a Resource | 398

 | Delete a Resource Dependency | 399

Filter Resources | 399

Graphs | 401

Charts Overview | 401

Manage the Charts Page | 404

Monitor Network Devices from the Charts Page | 409

 | Monitor Interface Flaps for a Single Interface | 409

 | Monitor Interface Flaps for 'ge' Interfaces on a Device | 411

 | Monitor System Memory Usage of all Devices in a Device Group | 413

 | Monitor RE CPU Usage for Multiple Devices | 416

 | Monitor RE CPU Usage for Multiple Devices side-by-side | 419

Grafana | 423

Grafana Overview | 423

Monitor Network Device Health Using Grafana | 424

 | Run a Query | 424

 | View Pre-populated Graphs | 426

7

Manage Sensor Settings, Insights Settings, and Data Summarization Profiles

Sensor Settings | 428

Sensors Overview | 429

 | Native GPB | 431

 | NetFlow | 431

sFlow | 433

OpenConfig | 434

- gNMI-Encoded OpenConfig RPC | 434

- Device Configuration for OpenConfig | 435

Syslog | 438

Server Monitoring Sensor | 440

iAgent (CLI/NETCONF) | 446

- Define PyEZ Table/View | 447

- Gather Output from Device | 448

- Generate JSON for Use in Paragon Automation Database | 449

- Outbound SSH (Device-Initiated) | 451

SNMP | 453

About the Ingest Settings Page | 454

Configure NetFlow Settings | 455

- Use Pre-defined NetFlow Templates | 456

- Create Custom NetFlow Templates | 456

- Delete a NetFlow Template | 457

- Clone an Existing NetFlow Template | 458

- Configure Flow Source IP Address | 459

- Configure Flow Ports | 460

Configure a Rule Using Flow Sensor | 461

About the Frequency Profiles | 468

Manage Frequency Profiles | 469

- Configure a Frequency Profile | 470

- Edit a Frequency Profile | 470

- Clone a Frequency Profile | 471

- Delete a Frequency Profile | 471

Apply a Frequency Profile | 473

Configure Offset Time | 474

- Offset Used in Formulae | 475

- Offset Used in Reference | 476

- Offset Used in Vectors | 477

- Offset Used in Trigger Term | 478

- Offset Used in Frequency Profile Applied to a Rule | 479
- Configure a Rule Using Server Monitoring Sensor | 482
- Configure Native GPB Ingest | 485
- Configure sFlow Settings | 486
 - Configure Devices to Send sFlow Packets | 487
 - Configure sFlow Ingest | 488
 - Delete sFlow Settings | 493
 - Configure sFlow in Devices and Device Groups | 496
 - Configure a Rule Using sFlow | 498
- Configure SNMP Ingest | 499
- Configure a Rule Using SNMP Scalar | 503
- Configure SNMP Trap and Inform Notifications | 504
 - Tasks You Can Perform | 504
 - Find the Engine ID | 506
 - Configure Trap Notifications | 506
 - Configure Inform Notifications | 511
 - Configure Port for Inform Notifications | 512
 - Configure a Rule for SNMP Notification | 513
- Configure Outbound SSH Port for iAgent | 515
- Configure System Log Ingest | 516
 - Device Configuration | 517
 - Configure Syslog Events Pattern | 517
 - Add Patterns to a Pattern Set | 520
 - Configure Header Pattern | 520
 - Edit a Header Pattern | 523
 - Clone a Syslog Events Pattern | 523
 - Clone a Pattern Set | 524
 - Configure Multiple Source IP Addresses for a Device | 524
- System Log Optional Configurations | 525
 - Configure Syslog Ports | 525
 - Configure Syslog Time Zone | 525
 - Configure Host Name Aliases for a Device | 526

- Configure a Rule Using Syslog | 526
- Understand Inband Flow Analyzer 2.0 | 532
- Configure Device Details for Inband Flow Analyzer Devices | 538
- Delete an Inband Flow Analyzer Device | 539
- Understand Bring Your Own Ingest | 540
- Load BYOI Default Plug-ins | 541
- Configure Bring Your Own Ingest Default Plug-in Instances | 542
- Build and Load BYOI Custom Plug-in Images | 544
 - Create a Process File for the Plug-in Image | 545
 - Create a Shell Script for Configuration Updates | 548
 - Tag and Export the BYOI Custom Plug-in Image | 548
 - Configure Kubernetes YAML File | 549
 - Assign Virtual IP Address to Plug-in | 552
 - Load the BYOI Custom Plug-in | 554
- Configure Bring Your Own Ingest Custom Plug-in Instances | 555
- Use Sample Rule and Playbook Configurations for BYOI Custom Plug-in Instances | 557
- Configure Ingest Mapping for Default BYOI Plug-in Instances | 558
- Delete a BYOI Plug-in | 560
- About the Diagnostics Page | 561
- Use the Self Test Tool | 564
- Use the Reachability Test | 566
- Use the Ingest Test Tool | 567
- Use the No-Data Tool | 568
- Paragon Insights Tagging Overview | 570
- Types of Tagging | 577
- Add a Tagging Profile | 585
- Apply a Tagging Profile | 590

Delete a Tagging Profile | 591

Understand User-Defined Actions and Functions | 593

Modify User-Defined Action, Function, and Workflow Engines | 594

Enable UDA Scheduler in Trigger Action | 599

Understand kube-state-metrics Service | 600

Insights Settings | 613

About the Insights Settings Page | 613

Add Alert Blackouts | 617

About Alert Notifications | 619

Use Exim4 for E-mails | 619

Configure the Exim4 Agent to Send E-mail | 620

Configure a Notification Profile | 621

Enable Alert Notifications for Device Groups and Network Groups | 627

Configure Report Settings | 628

Configure Scheduler Settings | 630

Configure a Retention Policy | 632

Configure Destination Settings | 633

Time Series Database (TSDB) Overview | 635

Manage Time Series Database Settings | 638

Backup and Restore the TSDB | 640

Time Series Database Replication Scenarios | 642

Points to Remember | 642

Scenario One | 643

Scenario Two | 644

Scenario Three | 645

Frequently Asked Questions | 645

Data Summarization Profiles | 647

Data Summarization Overview | 647

About the Raw Data Summarization Profiles Page | 649

About the Data Roll Up Summarization Profiles Page | 650

Add a Raw Data Summarization Profile | 651

Add a Data Rollup Summarization Profile | 653

Apply Data Summarization Profiles | 655

8

Configure Your Network

Topology | 660

Interactive Map Features Overview | 660

About the Topology Page | 677

Left Widget Options on Topology Page | 680

Group Nodes | 682

Group Nodes and Links into a Topology Group | 683

Ungroup Nodes | 684

Automatically Group Nodes | 684

Manage Map Layouts | 686

Network Information Table | 688

Network Information Table Overview | 689

About the Node Tab | 692

Add a Node | 696

Edit Node Parameters | 699

Delete a Node | 701

About the Link Tab | 702

Add a Link | 705

Edit Link Parameters | 708

Delete a Link | 709

About the Tunnel Tab | 710

Understand How Pathfinder Handles LSPs | **715**

Reroute LSPs Overview | **718**

Segment Routing Overview | **719**

Add a Single Tunnel | **729**

Add Diverse Tunnels | **743**

Add Multiple Tunnels | **754**

Edit and Delete Tunnels | **764**

 | Edit Tunnels | **764**

 | Delete Tunnels | **765**

About the Demand Tab | **765**

About the Interface Tab | **767**

Container LSP Overview | **768**

About the Container LSP Tab | **769**

Add a Container LSP | **770**

Edit Container LSP Parameters | **775**

Maintenance Event Overview | **775**

About the Maintenance Tab | **777**

Add a Maintenance Event | **778**

Edit a Maintenance Event | **780**

Simulate a Maintenance Event | **781**

Delete a Maintenance Event | **782**

About the P2MP Groups Tab | **783**

Add a P2MP Group | **791**

Edit P2MP Group Parameters | **798**

About the SRLG/Facility Tab | **799**

Add an SRLG/Facility | **800**

Edit SRLG/Facility Parameters | 801

About the Topology Group Tab | 801

Add Anycast Group Tunnels | 803

Tunnels | 815

Understand LSP Delegation and Undelegation | 815

Add and Remove LSP Delegation | 816

 Add LSP Delegation | 817

 Remove LSP Delegation | 818

About the Bandwidth Calendar Page | 819

About the Path Optimization Page | 821

Change Control Management | 824

Change Control Management Overview | 824

Change Request Workflow | 826

About the Change Control Management Page | 829

9

Monitoring

Monitor Network Health | 833

About the Network Health Page | 833

Activate Time Inspector View | 844

Manage Alarms and Alerts | 845

Alerts and Alarms Overview | 845

About the Alarms Page | 847

About the Alerts Page | 850

Monitor Jobs | 854

About the Jobs Page | 854

Viewing Job Details | 856

View Job Status | 856

Analytics | 858

Configure Routers to Send JTI Telemetry Data and RPM Statistics to the Data Collector | 858

Configure Routers to Advertise Link Statistics through BGP-LS | 861

NetFlow Collector Overview | 864

Collect Analytics Data Overview | 873

View Analytics Data | 882

Monitor Workflows | 899

Action Engine Workflow Overview | 899

About the Workflows Monitor Page | 900

About the Workflows Page | 902

Manage Action Engine Workflows | 903

Add an Action Engine Workflow | 903

Run an Action Engine Workflow | 907

Stop an Instance | 908

Resume a Suspended Instance | 908

Filter Instances | 909

Delete an Action Engine Workflow | 909

10

Reports

Health Reports | 911

About the Health Reports Page | 911

View and Download Health Reports | 912

Compare Differences in Health Reports | 913

Network Reports | 914

Network Reports Overview | 914

View Network Reports | 915

Maintenance Reports | 918

Maintenance Reports Overview | 918

View Maintenance Reports | 919

Inventory Reports | 934

Inventory Reports Overview | 934

View Inventory Reports | 936

Demand Reports | 947

Demand Reports Overview | 947

View Demand Reports | 948

11

Administration

Manage E-mail Templates | 953

E-mail Templates Overview | 953

About the E-mail Templates Page | 954

Edit an E-mail Template | 955

Manage Audit Logs | 957

Audit Logs Overview | 957

About the Audit Logs Page | 959

Filter Audit Logs | 960

Export Audit Logs | 961

Configure External EMS | 963

About the External EMS Page | 963

Add an External EMS | 965

Edit and Delete an External EMS | 967

 | Edit an External EMS | 967

 | Delete an External EMS | 967

Manage Task Scheduler | 969

About the Task Scheduler Page | 969

Add a Bandwidth Sizing Task | 972

Add a Container Normalization Task | 975

Add a Device Collection Task | 978

Add a Demand Aging Task | 983

[Add a Demand Reports Task | 985](#)

[Add a Network Archive Task | 990](#)

[Add a Network Maintenance Task | 993](#)

[Add a Network Cleanup Task | 997](#)

[Edit and Delete Tasks | 999](#)

[Edit Tasks | 1000](#)

[Delete Tasks | 1000](#)

[Manage Security Settings | 1001](#)

[About the Security Settings Page | 1001](#)

[Configure Security Profiles for SSL Authentication | 1003](#)

[License Management | 1005](#)

[Paragon Insights Licensing Overview | 1005](#)

[About the License Management Page | 1006](#)

[View, Add, or Delete Licenses | 1009](#)

About This Guide

Use this guide to understand the features and tasks that you can configure and perform from the Paragon Automation GUI. This guide provides feature overviews and procedures that help you understand the features and perform Paragon Automation configuration tasks.

1

PART

Introduction

[Overview](#) | 2

[Access the Paragon Automation GUI](#) | 35

[Access the Paragon Planner](#) | 48

[Configure SMTP, LDAP, and Portal Settings](#) | 50

[Manage Users](#) | 57

[Manage Roles](#) | 66

[Manage User Groups](#) | 74

[Identity Providers](#) | 81

CHAPTER 1

Overview

IN THIS CHAPTER

- [About the Paragon Automation User Guide | 2](#)
- [Paragon Automation Overview | 3](#)
- [Paragon Pathfinder Overview | 6](#)
- [Paragon Insights Overview | 9](#)
- [Paragon Planner Overview | 13](#)
- [Understand Differences between Paragon Pathfinder and Planner | 15](#)
- [Paragon Automation GUI Overview | 18](#)
- [Paragon Automation GUI Menu Overview | 27](#)

About the Paragon Automation User Guide

This guide provides an overview of Paragon Automation and its applications, and describes how to use the Paragon Automation GUI to perform tasks using the different Paragon Automation applications.

This guide is available on the [Paragon Automation Documentation page](#), which contains links to the different Paragon Automation releases and the documentation published for those releases. [Table 1 on page 2](#) lists some of the other documentation that is available on the Paragon Automation Documentation page.

Table 1: Additional Paragon Automation Documentation

Documentation	Location
Paragon Automation Release Notes	DISCOVER section of the Paragon Automation Documentation page

Table 1: Additional Paragon Automation Documentation (Continued)

Documentation	Location
Paragon Automation Getting Started Guide	SET UP section of the Paragon Automation Documentation page
Paragon Automation Installation Guide	SET UP section of the Paragon Automation Documentation page
Paragon Planner Desktop Application User Guide	MANAGE section of the Paragon Automation Documentation page
Paragon Planner User Guide	MANAGE section of the Paragon Automation Documentation page
Paragon Automation Troubleshooting Guide	HELP RESOURCES > GUIDES section of the Paragon Automation Documentation page

NOTE: Paragon Active Assurance is not yet integrated with the Paragon Automation suite of applications. So, you can access the documentation from the [Paragon Active Assurance](#) technical documentation page.

RELATED DOCUMENTATION

| [Paragon Automation Overview](#) | 3

Paragon Automation Overview

IN THIS SECTION

- [Key Use Cases](#) | 5
- [Benefits of Paragon Automation](#) | 5

As networks have become pervasive and more important in our daily lives, the stakes have increased for ensuring high-quality service and operational experience. In addition, networks have expanded in scale, making them extremely difficult to manage manually. Because human error causes many network outages, an increase in scale means a greater likelihood of human error.

To improve the operational and service experience, it's essential for network operators to reduce the number of network errors, improve resolution times, and gain better visibility into and control of their networks. Automation is the key to realizing these outcomes.

Paragon Automation by Juniper Networks is a modular portfolio of cloud-native software applications that help you simplify network operations by eliminating manual tasks, processes, and workflows that are often repetitive and prone to human error. Paragon Automation delivers closed-loop automation to translate business intent into service performance across the entire service-delivery lifecycle. Paragon Automation builds on Juniper's existing automation portfolio to meet the most pressing challenges of current and next-generation networks and services.

The Paragon Automation portfolio includes the following products, which enable you to perform the network tasks shown in the list:

- Paragon Pathfinder (formerly NorthStar Controller): Design, provision, and optimize segment routing (SR) and MPLS path flows. See ["Paragon Pathfinder Overview" on page 6](#).
- Paragon Planner (formerly NorthStar Planner): Plan, model, and verify services before deploying them; forecast the impact of network changes such as latency, traffic flows, and new traffic; and forecast the impact of new services. See ["Paragon Planner Overview" on page 13](#).
- Paragon Active Assurance (formerly Netrounds): Actively test and assure services across physical, hybrid, and virtual networks. Paragon Active Assurance utilizes synthetic traffic to verify application and service performance at the time of service delivery and throughout the life of the service. See the [Paragon Active Assurance](#) technical documentation page.

NOTE: Currently, Paragon Active Assurance features are not accessible from the Paragon Automation GUI because Paragon Active Assurance is not yet integrated with the Paragon Automation suite of applications.

- Paragon Insights (formerly HealthBot): Diagnose network health and services health using streaming telemetry and machine learning (ML) analytics, which provide actionable insights into network behavior to help you identify and resolve issues quickly. See ["Paragon Insights Overview" on page 9](#).
- Anuta Networks ATOM: Provides multi-vendor provisioning, orchestration, and workflow automation for networks and services to eliminate costly errors while increasing operational speed, accuracy, and quality. See the [Anuta ATOM Resource Center](#) page.

Key Use Cases

- Automated service provisioning and monitoring: Service deployment with automated service monitoring.
- Software-defined networking (SDN) IP transport management: Automated management of IP transport using an SDN controller.
- Zero-touch testing: Automatically verify service quality prior to onboarding customers.
- Automated root cause analysis: Spot a service quality problem and root cause and remedy the problem.
- Anomaly detection: Understand network anomalies and assess impact.
- Coordinated maintenance: Fix issues without traffic impact.

Benefits of Paragon Automation

- Shorten time to revenue for new infrastructure and services by leveraging the power of automation.
- Assure service experience throughout the service lifecycle by using active testing, monitoring, predictive analytics, and closed-loop remediation.
- Minimize risk associated with upgrades, changes, and new deployments, and reduce failed service delivery rates, mean-time-to-repair (MTTR), and SLA penalties. Minimizing these risks results in higher service uptime, better quality of service, and increased customer satisfaction and retention.
- Obtain OPEX savings, higher workforce efficiency, and increased network utilization by automating device and service deployment, management, and compliance. You realize these improvements by using workflow automation for greater network visibility and repeatable, best-in-class service delivery.

RELATED DOCUMENTATION

| [Paragon Automation GUI Overview](#) | 18

Paragon Pathfinder Overview

IN THIS SECTION

- [Key Features | 7](#)
- [Key Use Cases | 7](#)
- [Benefits of Paragon Pathfinder | 8](#)

As more content and applications are migrated to the cloud and as new services are delivered over the network, network operators are expanding their networks to meet the increase in bandwidth demand. This expansion and the growing number of features in the network are increasing the complexity of traffic management. Therefore, network operators need to find a way to manage this complexity and to deliver increased speed and agility in their networks.

The high bandwidth demand and the increase in the number of latency-sensitive applications also means that network operators must meet stringent service-level agreements (SLAs). Further, to ensure that they reduce costs (CapEx and OpEx) and maximize revenue, network operators must run their networks hotter and more efficiently to make greater use of network bandwidth, potentially eliminating the need for redundant paths. Network operators must also design, implement, and operate their networks to make them operationally efficient. To achieve these outcomes, network operators are using automation.

Software-defined networking (SDN) controllers help network operators to visualize, monitor, and automate their network by using closed-loop automation. In addition, segment routing (SR) simplifies traffic management of IP-MPLS networks while integrating application awareness into the network control plane, thus providing the best possible application quality of experience (AppQoE) without increasing network complexity.

Juniper's Paragon Pathfinder (formerly NorthStar Controller) is a standards-based, stateful SDN and segment routing controller, which enables granular visibility and control of IP-MPLS and SR traffic flows in large service provider, cloud provider, and enterprise networks. Paragon Pathfinder collects topology and performance statistics, and provides network operators with a view of the network through real-time topology view and events, traffic and latency graphs, and traffic reports.

[Table 2 on page 7](#) displays the key functions of Paragon Pathfinder.

Table 2: Key Functions of Paragon Pathfinder

Category	Description
Analyze	Discovers the network topology by gathering data using routing protocols such as BGP and OSPF.
Optimize	Computes the path for services by using topology and user-defined constraints and by analyzing data to take intelligent decisions, thereby ensuring that applications follow the most efficient path across the network and meet SLA requirements.
Deploy	Installs the path by using protocols such as Path Computation Element Protocol (PCEP), BGP Segment Routing-Traffic Engineering (SR-TE), NETCONF, or YANG.

Key Features

- Provides real-time topology view and view of events in the network by using BGP Link State (BGP-LS), PCEP, gRPC, and NETCONF.
- Enables centralized discovery and provisioning, monitoring, and management of label-switched paths (LSPs).
- Enables complex, inter-domain path computation by using sophisticated algorithms.
- Facilitates the archival of network data, which can then be used for network planning (by using Paragon Planner).
- Integrates with Paragon Insights to monitor and manage device and network health.
- Provides open southbound (for example, PCEP, BGP-LS, NETCONF) and northbound interfaces (for example, REST, YANG)

Key Use Cases

- Visualization: Visualize and monitor MPLS networks running RSVP-Traffic Engineering (RSVP-TE), LDP, or SR-TE.
- Path diversity: Use traffic engineering to achieve a diverse path through the network. For example, if you have a critical service and want to guarantee SLAs on that service, you can specify a diverse path through the network, so that if one path fails, the traffic can move to a different path as quickly as possible.
- Rerouting LSPs based on different criteria:

- End-to-end utilization threshold violation: Paragon Pathfinder tracks the threshold for each interface. If the thresholds are violated (compared to user-defined global threshold values), the LSPs are rerouted based on priority, bandwidth, and so on.
- Delay threshold violation: Paragon Pathfinder collects the measured delay and reroutes LSPs that are transiting on links that violate a configured maximum delay.
- Packet loss threshold violation: Paragon Pathfinder collects packet loss data and compares it with the user-defined global threshold values. Any link that violates the packet loss threshold is placed in maintenance mode for an hour, and LSPs that transit the link are rerouted based on priority, bandwidth, and so on.
- Node maintenance: If a user puts a node in maintenance (thereby triggering a maintenance event), Paragon Pathfinder excludes all nodes placed under maintenance from path computations, and automatically reroutes the affected LSPs.
- Bandwidth calendaring: Schedule future bandwidth needs.

Benefits of Paragon Pathfinder

- Provides better service experience by enabling network operators to monitor their networks for real-time topology and bandwidth changes, and optimizes network services to deliver high-quality customer experience.
- Simplifies the operational experience because network operators can see an overview of their entire MPLS or segment routed networks and then drill down to view the detailed state of a path, link, or node.
- Enables the network to operate more autonomously by enabling MPLS and segment-routed networks to self-learn changes in topology, bandwidth usage, and traffic patterns and then to take appropriate action to maintain SLAs.

RELATED DOCUMENTATION

| [Paragon Planner Overview](#) | 13

Paragon Insights Overview

IN THIS SECTION

- Main Components of Paragon Insights | 9
- Closed-Loop Automation | 11
- Benefits of Paragon Insights | 13

Paragon Insights is a highly automated and programmable device-level diagnostics and network analytics tool that provides consistent and coherent operational intelligence across network deployments. Paragon Insights integrates multiple data collection methods (such as Junos Telemetry Interface (JTI), NETCONF, system log [syslog], and SNMP), to aggregate and correlate large volumes of time-sensitive telemetry data, thereby providing a multidimensional and predictive view of the network. Additionally, Paragon Insights translates troubleshooting, maintenance, and real-time analytics into an intuitive user experience to give network operators actionable insights into the health of individual devices and of the overall network.

Main Components of Paragon Insights

Paragon Insights consists of two main components:

- Health Monitoring, which enables you to:
 - View an abstracted, hierarchical representation of device and network-level health.
 - Define the health parameters of key network elements through customizable key performance indicators (KPIs), rules, and playbooks.

A playbook is a collection of rules. You can create a playbook and apply the playbook to a device group or a network group. For more information on rules and playbooks, see [Paragon Insights Rules and Playbooks](#).

- Root Cause Analysis (RCA), which helps you find the root cause of a device or network-level issue when Paragon Insights detects a problem with a network element.

NOTE: The RCA feature is disabled, by default, in Paragon Automation Release 23.1. For instructions on how to enable the feature, see the [Release Notes](#).

Paragon Insights Health Monitoring

The Challenge

With increasing data traffic generated by cloud-native applications and emerging technologies, service providers and enterprises need a network analytics solution to analyze volumes of telemetry data, offer insights into overall network health, and produce actionable intelligence. Although telemetry-based techniques have existed for years, the growing number of protocols, data formats, and key performance indicators (KPIs) from diverse networking devices has made data analysis complex and costly. Traditional CLI-based interfaces require specialized skills to extract business value from telemetry data, creating a barrier to entry for network analytics.

How Paragon Insights Health Monitoring Helps

By aggregating and correlating raw telemetry data from multiple sources, the Paragon Insights Health Monitoring component provides a multidimensional view of network health that reports current status and projected threats to the network infrastructure and its workloads.

Health status determination is tightly integrated with the Paragon Insights RCA component, which can use system log data received from the network and its devices. Health Monitoring provides status indicators that alert you when a network resource is currently operating outside a user-defined performance policy. Health Monitoring does a risk analysis using historical trends and predicts whether a resource may become unhealthy in the future. Health Monitoring not only offers a fully customizable view of the current health of network elements, but also automatically initiates remedial actions based on predefined service level agreements (SLAs).

Defining the health of a network element, such as broadband network gateway (BNG), provider edge (PE), core, and leaf-spine, is highly contextual. Each element plays a different role in a network, with unique KPIs to monitor. Because there's no single definition for network health across all use cases, Paragon Insights provides a highly customizable framework to allow you to define your own health profiles.

Paragon Insights Root Cause Analysis

The Challenge

For some network issues, it can be challenging for network operators to determine what caused a networking device to stop working properly. In such cases, an operator must consult a specialist (with knowledge built from years of experience) to troubleshoot the problem and find the root cause.

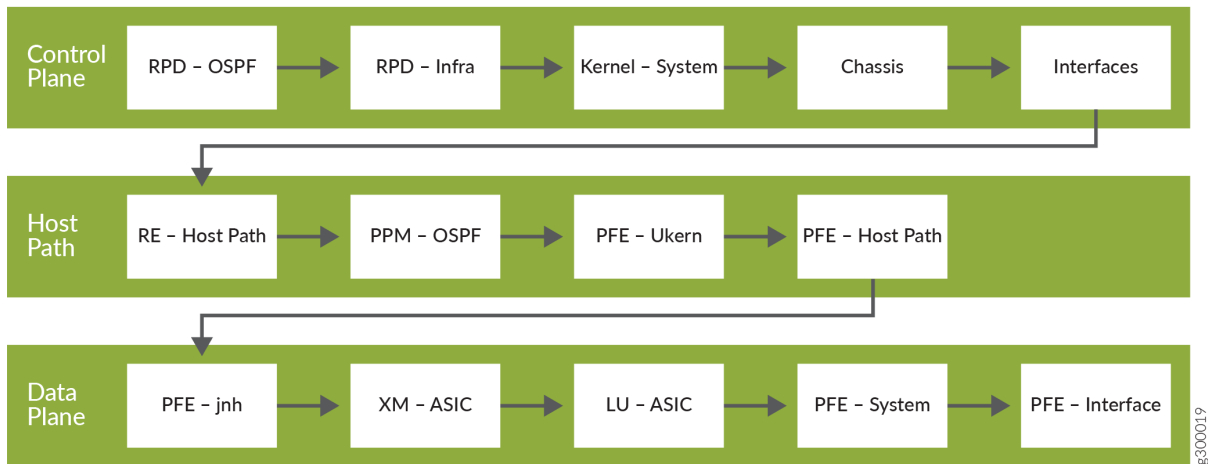
How Paragon Insights RCA Helps

The Paragon Insights RCA component simplifies the process of finding the root cause of a network issue. RCA captures the troubleshooting knowledge of specialists and has a knowledge base in the form of Paragon Insights rules. These rules are evaluated either on demand by a specific trigger or periodically in

the background to ascertain the health of a networking component (such as routing protocol, system, interface, or chassis) on the device.

To illustrate the benefits of Paragon Insights RCA, let's consider the problem of OSPF flapping. [Figure 1 on page 11](#) highlights the workflow sequence involved in debugging OSPF flapping. A network operator troubleshooting this issue would need to perform manual debugging steps for each tile (step) of the workflow sequence in order to find the root cause of the OSPF flapping. On the other hand, the RCA component troubleshoots the issue automatically by using an RCA bot. The RCA bot tracks all of the telemetry data collected by Paragon Insights and translates the information into graphical status indicators (displayed in the Paragon Insights web GUI) that correlate to different parts of the workflow sequence shown in [Figure 1 on page 11](#).

Figure 1: High-level workflow to debug OSPF flapping



When you configure Paragon Insights, each tile of the workflow sequence (shown in [Figure 1 on page 11](#)) can be defined by one or more rules. For example, the RPD-OSPF tile could be defined as two rule conditions: one to check if "hello-transmitted" counters are incrementing and the other to check if "hello-received" counters are incrementing. Based on these user-defined rules, Insights provides status indicators, alarm notifications, and an alarm management tool to inform and alert you of specific network conditions that could lead to OSPF flapping.

By isolating a problem area in the workflow, RCA proactively guides you in determining the appropriate corrective action to take to fix a pending issue or avoid a potential one.

Closed-Loop Automation

Paragon Insights offers closed-loop automation. The automation workflow can be divided into seven main steps (see [Figure 2 on page 12](#)):

Benefits of Paragon Insights

- Customization—Provides a framework to define and customize health profiles, allowing truly actionable insights for the specific device or network being monitored.
- Automation—Automates root cause analysis and log file analysis, streamlines diagnostic workflows, and provides self-healing and remediation capabilities.
- Greater network visibility—Provides advanced multidimensional analytics across network elements, giving you a clearer understanding of network behavior to establish operational benchmarks, improve resource planning, and minimize service downtime.
- Intuitive GUI—Offers an intuitive web-based GUI for policy management and easy data consumption.
- Open integration—Lowers the barrier of entry for telemetry and analytics by providing open source data pipelines, notification capabilities, and third-party device support.
- Multiple data collection methods—Includes support for JTI, OpenConfig, NETCONF, CLI, Syslog, NetFlow, and SNMP.

RELATED DOCUMENTATION

| [Paragon Insights Getting Started Guide](#)

Paragon Planner Overview

IN THIS SECTION

- [Key Features and Use Cases | 14](#)
- [Benefits of Paragon Planner | 15](#)

Network automation has become increasingly important to operators who want to improve their operational and service experience. Network operators can apply automation across the entire network life cycle: from planning and designing, to implementing and operating, and optimizing the network. In addition, by using feedback from the different steps in the lifecycle, operators can achieve closed loop network automation.

The first step in realizing closed loop network automation is the accurate planning and designing of the network. Juniper's Paragon Automation suite of applications includes Paragon Planner, which helps operators to plan and design their network.

Paragon Planner is a network modeling tool that enables offline visualization of network resources and provides detailed architectural planning of production networks. Network operators can use Paragon Planner to visualize and map services to resources, and forecast the impact of network changes (such as latency, additional traffic, shifts in traffic flows, and new capacity) on transport services. Paragon Planner also enables operators to simulate network changes and other traffic scenarios without affecting the production network, and assess the network for potential failure scenarios.

Key Features and Use Cases

The key features of and use cases for Paragon Planner include the following:

- Automatically construct network topologies by using data snapshots of live networks (obtained from Paragon Pathfinder) or source data from stored network configuration files and other sources.
- Perform capacity planning to determine whether there is sufficient capacity or if more capacity should be added, and which links can be pruned without compromising resiliency.
- Analyze traffic loads to determine accurate link utilization for failure simulation.
- Validate network changes (in a safe, virtual environment) before deployment.
- Run what-if scenarios to anticipate the impact of any network change.
- Create and model VPNs, simulate VPN routing, and generate VPN traffic.
- Simulate multicast flows based on user-defined multicast groups and demands.
- Design and simulate MPLS traffic engineering (MPLS-TE) and label-switched path (LSP) routing.
- Assess network resiliency against different failure scenarios and analyze how traffic is rerouted and its effect on network links.
- Ensure service-level agreement (SLA) compliance by modeling class of service (CoS) classes and policies and queuing schemes. (You can define application flows based on CoS to enable the modeling of voice over IP (VoIP) or video on demand (VOD) traffic.)
- Model and analyze BGP routing.
- Hardware inventory.
- Network integrity checks.

Benefits of Paragon Planner

- Lower CapEx and OpEx (hardware and maintenance costs) by using Planner's tariff-based design, MPLS LSP, and segment routing traffic engineering (SR-TE) features for effective utilization of the network, and superior design optimization.
- Quickly diagnose performance problems by using flow analysis, bottleneck detection and analysis, peak utilization analysis, and multicast simulation features.
- Optimize plans for future network growth (to meet business needs) by using capacity planning and data forecasting.
- Avoid problems and mitigate risk by assessing the network using Planner's resiliency analysis, fiber cut analysis, and so on.
- Validate new services, equipment, and technologies before they are rolled out.

RELATED DOCUMENTATION

[Paragon Pathfinder Overview | 6](#)

[Access Paragon Planner Desktop Application | 48](#)

Understand Differences between Paragon Pathfinder and Planner

This topic explains key differences between Paragon Pathfinder and Paragon Planner, two applications that belong to the Paragon Automation suite of applications:

- Paragon Pathfinder (formerly NorthStar Controller) is a traffic engineering solution that simplifies and automates provisioning, management, and monitoring of segment routing and IP/MPLS flows across large networks.
- Paragon Planner (formerly NorthStar Planner) is a network modeling tool that can be used for offline visualization and detailed architectural planning of any production network.

[Table 3 on page 16](#) lists the key differences between the two applications.

Table 3: Differences between Paragon Pathfinder and Paragon Planner

Paragon Pathfinder	Paragon Planner
Paragon Pathfinder enables you to monitor your live network and any changes that you make in Pathfinder are propagated to your live network.	Paragon Planner is an offline modeling application, which means that changes that you make in Planner do not affect your live network.
Pathfinder can connect to the live network because of connectivity between Pathfinder and the live network.	Planner has no capability to connect to the network because there's no connectivity between Planner and the network.
In Pathfinder, there's only one live network model.	In Planner, you can have one or more offline network models, which means you can analyze and compare impact of different design changes or failure scenarios. However, the changes are not propagated into the live network.
<p>In Pathfinder, the network model is based on the traffic-engineering database, which is based on the live (real time) status of the network.</p> <p>For example, if you configure a link that later goes down, the down status of the link is reflected in the traffic-engineering database, but not in the router's configuration.</p>	<p>The network models that we use in Planner are from archives and collections in Pathfinder, which means that Planner relies on data from Pathfinder.</p> <p>Planner builds the model of the network using router configuration files, which means that it uses a configuration-based model of the network. Therefore, Planner displays an intent-based model (configuration) with some supplemental live information from Pathfinder, such as initiated label-switched paths (LSPs), which are not available in the router's configuration.</p>
In Pathfinder, any changes you make to the network model affects the live network.	Planner allows you to run what-if scenarios, which allow you to make changes to the network model and see the effect of those changes, without affecting the live network.
In Pathfinder, you can perform a maintenance event simulation. You define a maintenance event by specifying what happens, for example, a router goes down or a link fails. Then, based on that event, you simulate what happens to the live LSPs and traffic demands.	In Planner, you can run failure simulations (for example, fail a link or a router) and analyze the impact on traffic demand or LSP tunnels.

Table 3: Differences between Paragon Pathfinder and Paragon Planner (Continued)

Paragon Pathfinder	Paragon Planner
The network topology map shows live node status, link utilization, and LSP paths.	The network topology map shows simulated or imported data for nodes, links, and LSP paths.
Network information table shows live status of nodes, links, and LSPs.	Network information table shows simulated or imported data for nodes, links, and LSPs.
Discover nodes, links, and LSPs from the live network by using Path Computation Element (PCE) protocol (PCEP), BGP-LS, or NETCONF.	Import and parse router configuration, or add nodes, links, and LSPs for network modeling.
Provision LSPs directly to the network.	Add and stage LSPs to the offline model for simulation. However, you cannot provision LSPs to the live network.
Create or schedule maintenance events to reroute LSPs around the impacted nodes and links.	Create or schedule simulation events to analyze the network model from failure scenarios.
Dashboard reports show the current status of and key performance indicators (KPIs) for the live network.	Report manager provides extensive reports for simulation and planning.
Collects real-time interface traffic or delay statistics, and stores the data for querying and for displaying in charts.	Import interface data or aggregate archived data to generate historical statistics for querying and displaying in charts.

RELATED DOCUMENTATION

[Paragon Pathfinder Overview | 6](#)
[Paragon Planner Overview | 13](#)

Paragon Automation GUI Overview

IN THIS SECTION

- [Menu Bar and Banner | 18](#)
- [Commonly Used Icons | 23](#)
- [Add or Remove Favorite Pages | 25](#)
- [User Inactivity | 26](#)

The Paragon Automation GUI provides an easy to use, single pane of glass experience that allows users to access and to use the different Paragon Automation applications.

After you log in successfully to the Paragon Automation GUI, the first page that you encounter is the Dashboard page, as shown in [Figure 3 on page 22](#). The Dashboard page displays multiple widgets (also known as dashlets) that you can customize. For more information, see ["About the Dashboard Page" on page 38](#).

In the rest of this topic, we'll discuss some commonly used elements and features of the Paragon Automation GUI.

Menu Bar and Banner

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

- **Menu bar:** The menu bar, which is available at the left-side of the Paragon Automation GUI, is minimized by default. You can mouse over or click inside the menu bar to expand the menu. A sample of the expanded menu is shown in [Figure 3 on page 22](#).

You can expand the high-level menu items and navigate to the different pages in the Paragon Automation GUI. For more information about the different menu entries and how they map to the different Paragon Automation applications, see ["Paragon Automation GUI Menu Overview" on page 27](#).

- **Banner:** The banner, which is displayed at the top of the page (see [Figure 3 on page 22](#)) contains several icons that you're likely to use regularly. The icons and their functions are explained in [Table 4 on page 19](#).

Table 4: Icons on the Paragon Automation Banner

Description	Function
Menu bar toggle	<p>Clicking the menu bar toggle icon (the icon with three horizontal bars) in the top left of the Paragon Automation banner to toggle the visibility of the Paragon Automation menu. This means that the menu is displayed if it was previously hidden and hidden if it was previously displayed.</p>
In-progress and scheduled jobs	<p>To quickly view the jobs that are running or that are scheduled, mouse over the clock icon, which opens a widget that contains two tabs: In-Progress and Scheduled.</p> <p>The number of jobs that are in progress and scheduled are shown in parentheses in the title of the tab, and each tab lists the in-progress and scheduled jobs.</p> <p>You can view all the jobs by clicking the See all jobs hyperlink, which takes you to the Jobs page.</p>

Table 4: Icons on the Paragon Automation Banner (*Continued*)

Description	Function
Deployment status	<p>Icons overlaid on top of the deployment status icon provide an easy way to find out the status of the last deployment job in Paragon Insights. These icons and their statuses are listed in Table 5 on page 21.</p> <p>When you mouse over the icon, a popup appears displaying the following:</p> <ul style="list-style-type: none"> • Last Result, which displays the status of and information related to the last configuration change. • Pending Deployments, which displays the number of affected services and the number of deleted services that occur because of the configuration changes. <p>If there are pending deployments, then you can commit or roll back the changes using the Health Configuration Deployment Status page, which you can access by clicking the deployment status icon. For more information, see "Commit or Roll Back Configuration Changes in Paragon Insights" on page 185.</p>
Insights playbooks and rules	<p>Click the Tools icon and select Rules (github) to access the Juniper Networks official and community supplied playbooks and rules for Paragon Insights.</p> <p>For details on how to add predefined rules and playbooks, see "Add a Predefined Rule" on page 340 and "Add a Predefined Playbook" on page 308.</p>
Feedback	<p>Click the Comment icon to send feedback (through e-mail) about your Paragon Automation instance.</p>

Table 4: Icons on the Paragon Automation Banner (Continued)

Description	Function
User menu	<p>Click the user (profile) icon to access the user menu. You can do the following:</p> <ul style="list-style-type: none"> • Change your password: Click Change Password to open the Change Password dialog and modify your password. See "Change Your Password" on page 45. • Log out of Paragon Automation: Click Logout to log out of the Paragon Automation GUI. <p>You are returned to the Paragon Automation login page.</p>
Help menu	<p>Click the Help (?) icon to access the help menu, which provides links to the following Paragon Automation documentation:</p> <ul style="list-style-type: none"> • Getting Started • What's New • Quick Help • Release Notes • About <p>You can also access the documentation on the Paragon Automation Documentation page.</p>

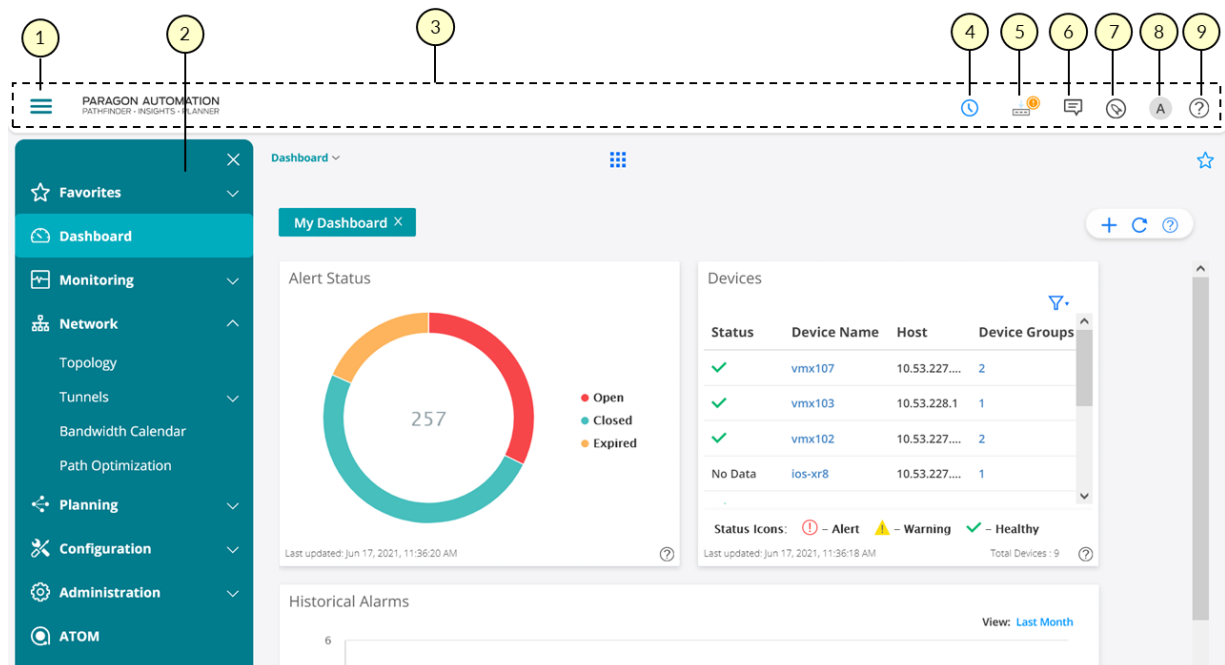
Table 5: Icons Indicating Last Deployment Status

Description of Icon	Meaning
Check mark (✓) in a green circle at the top-right of the Deployment status icon.	Last configuration change was successful.
Cross (×) in a red circle at the top-right of the Deployment status icon.	Last configuration change was unsuccessful.

Table 5: Icons Indicating Last Deployment Status (Continued)

Description of Icon	Meaning
Exclamation point (!) in an orange circle at the top-right of the Deployment status icon.	Last configuration change is in progress or the status is unknown.
No icon at the top-right of the Deployment status icon.	No last configuration change is present.
Bouncing, blue downward arrow (↓) on the Deployment Status icon.	One or more deployments are pending.

Figure 3: Sample Paragon Automation Dashboard Page Showing Menu and Banner



1– Menu toggle icon	6– Feedback icon
2– Expanded menu bar	7– Icon to access Insights playbooks and rules
3– Banner	8– User (profile) icon
4– Icon for viewing in-progress and scheduled jobs	9– Help (?) icon

5– Icon for viewing deployment status (Paragon Insights)

Commonly Used Icons

In this section, we list some of the common icons that you'll come across when you use the Paragon Automation GUI:

- [Figure 4 on page 24](#) shows the sort, filter, and other icons that you typically encounter on landing pages such as Alerts, Alarms, and Users. [Table 6 on page 23](#) provides a high-level explanation of their functions.

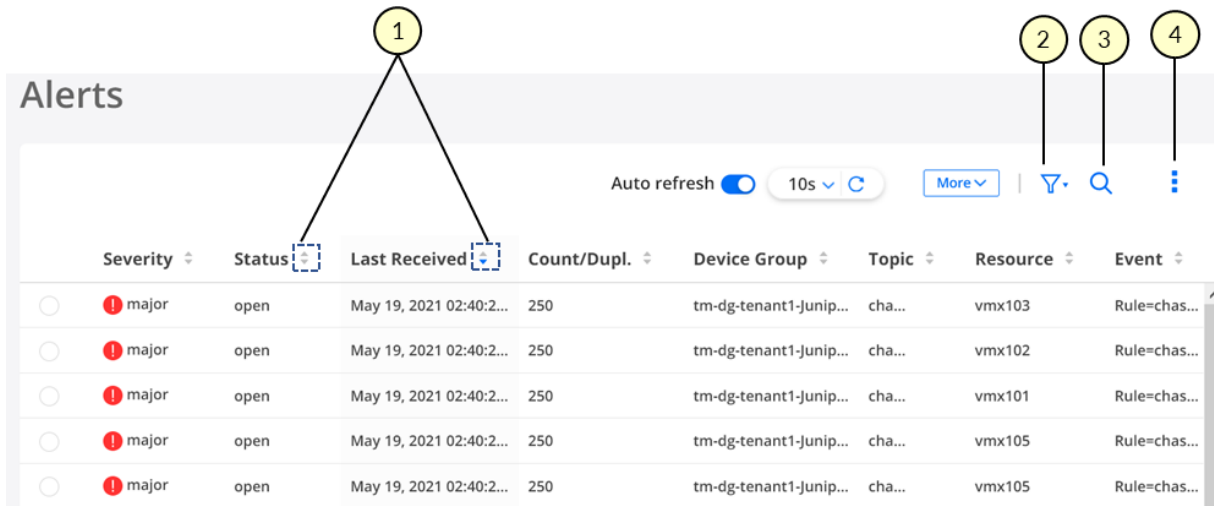
NOTE: The search and filter icons might not be available on some pages.

- [Figure 5 on page 25](#) shows the breadcrumbs, help, add, and other icons, and [Table 7 on page 24](#) provides a high-level explanation of their functions.

Table 6: Sort, Filter, Search, and Show or Hide Columns Icons

Description	Function
Sort status icons (up and down arrows)	Sort status 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 sort status icons next to the label indicates whether the data is sorted in ascending or descending order.
Filter icon (funnel)	Allows you to apply one or more filters to the data in the table and, if needed, save the filters.
Search icon (magnifying glass)	Allows you to search the data and, if needed, save the search as a filter.
Show or hide columns icon (three vertical dots or ellipses)	Allows you to pick the columns that you want displayed on the page or reset the preferences to the default.

Figure 4: Alerts Page Showing Sort, Filter, and Other Icons



1- Sort status icons	3- Search icon
2- Filter icon	4- Show or Hide columns icon

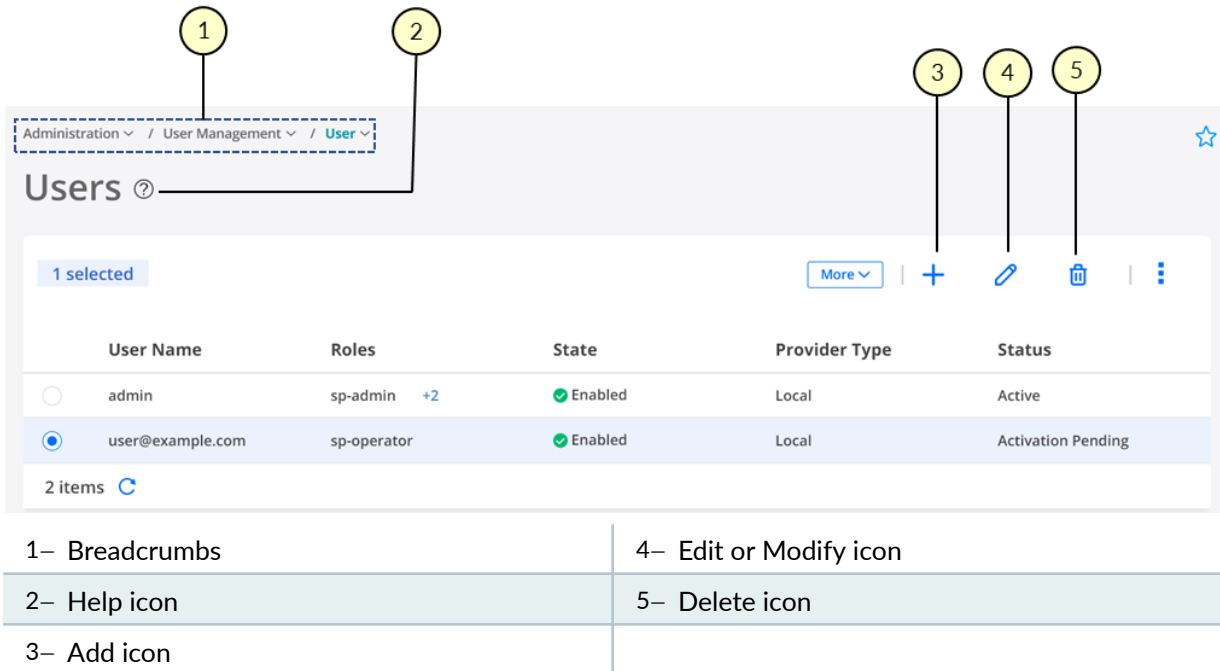
Table 7: Breadcrumbs, Help, and Other Icons

Name	Description
------	-------------

NOTE: Access to the add, modify, and delete capabilities depends on the role that you are assigned as a user. For more information, see ["Roles Overview"](#) on page 66.

Breadcrumbs	Provides an alternate way to navigate and access the Paragon Automation menu items.
Help icon (?)	Mouse over the help icon to get high-level information about the page that you're viewing.
Add icon (+)	Enables you to add objects; for example, add users or roles.
Modify icon (pencil)	Enables you to modify existing objects.
Delete icon (trash can)	Enables you to delete existing objects.

Figure 5: Breadcrumbs, Help, and Other Icons



Add or Remove Favorite Pages

The Favorites feature in the Paragon Automation GUI allows you to mark pages that you frequently use or visit as favorites, so that you can access such pages easily. [Figure 6 on page 26](#) shows a sample page with existing favorites.

- View or access favorite pages: You can use the Favorites sub-menu in the Paragon Automation menu to view the existing favorite pages. To access a favorite page, click on the title of the page (under Favorites).
- 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 a page title in the menu bar, as shown in [Figure 6 on page 26](#).
 - By clicking the star icon at the top right corner of a page (below the Paragon Automation banner), as shown in [Figure 6 on page 26](#).

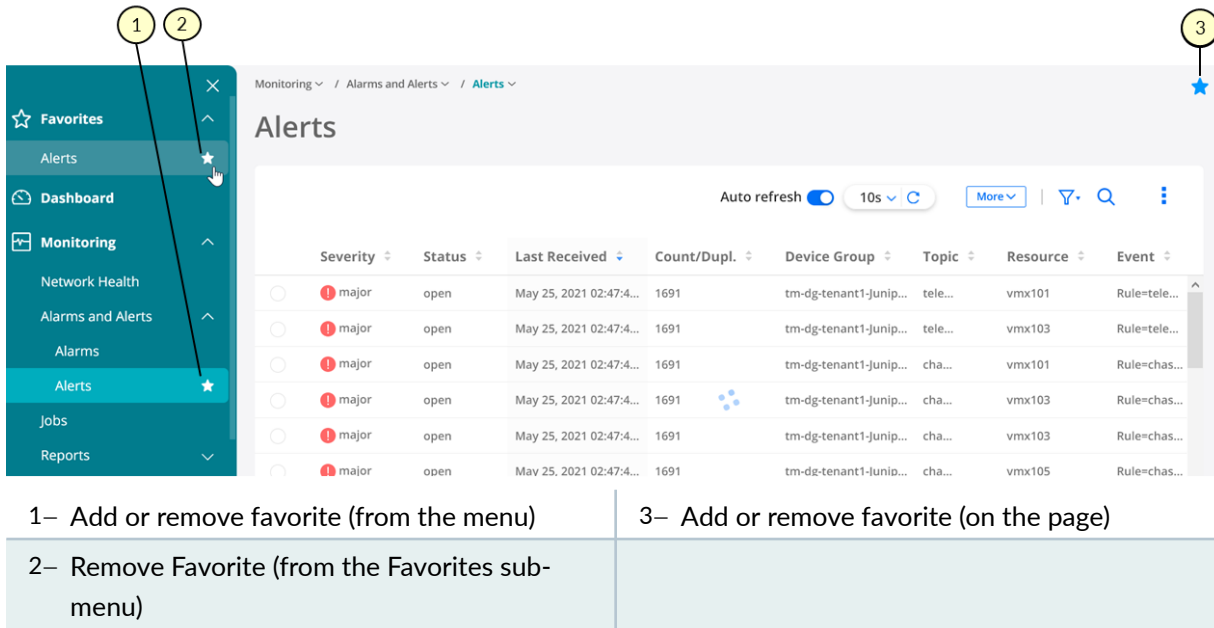
When you add a page as a favorite, it is displayed under the Favorites sub-menu, and the star icons are shaded (filled) indicating that the page is a favorite.

- Remove a page as a favorite: You can remove a page as a favorite in one of the following ways:
 - By clicking the star icon next to a page title under the Favorites sub-menu.
 - By clicking the star icon next to a page title in the menu bar.

- By clicking the star icon at the top right corner of a page.

When a page is removed as a favorite, it is removed from the Favorites sub-menu and the star icons change to empty (not shaded) indicating that the page is not a favorite.

Figure 6: View, Add, or Remove Favorites



User Inactivity

Any user who is idle (has not performed any keystrokes or mouse clicks) for 30 minutes is logged out of Paragon Automation. At the idle 25 minute mark, you will be prompted with a **Session Expiration** pop up where you can either **Extend** or **Log out** of the current session. From when the pop up is visible, you have five minutes to extend or log out of the current session. If you click **Extend**, the current session is extend by another 25 minutes.

RELATED DOCUMENTATION

[Access the Paragon Automation GUI | 36](#)

Paragon Automation GUI Menu Overview

IN THIS SECTION

- [Monitoring Sub-Menu | 28](#)
- [Reports Sub-Menu | 29](#)
- [Network Sub-Menu | 30](#)
- [Planning Sub-Menu | 31](#)
- [Configuration Sub-Menu | 31](#)
- [Administration Sub-Menu | 33](#)

The Paragon Automation GUI menu enables you to access the different Paragon Automation applications and perform tasks related to those applications. In addition, you can perform tasks that are common to the different Paragon Automation applications. The tasks that you can perform is based on the roles and access privileges (capabilities) that you're assigned as a Paragon Automation user. For more information, see ["Roles Overview" on page 66](#).

NOTE: Paragon Active Assurance is not yet integrated with the Paragon Automation suite of applications. Therefore, Paragon Active Assurance features are not accessible from the Paragon Automation GUI.

[Table 8 on page 27](#) shows the top-level menu items (sub-menus) in the Paragon Automation GUI.

Table 8: Paragon Automation Main Menu

Sub-menu	Description
Favorites	Displays the pages that are marked as favorites. See "Paragon Automation GUI Overview" on page 18 .
Dashboard	Access a user-configurable dashboard that you can customize with available widgets (also known as dashlets). See "About the Dashboard Page" on page 38 .

Table 8: Paragon Automation Main Menu (Continued)

Sub-menu	Description
Monitoring	Access various monitoring tasks, such as network health, alarms and alerts, and jobs. See Table 9 on page 29 .
Reports	Access various types of reports, such as health, network, maintenance, inventory, and demand reports. See Table 10 on page 29 .
Network	Access tasks related to Paragon Pathfinder, such as network topology, label-switched path (LSP) delegation, and path optimization. See Table 11 on page 31 .
Planning	Access the Paragon Planner desktop and web application. See Table 12 on page 31 .
Configuration	Access tasks related to configuration, such as adding and discovering devices, device and network groups, templates, playbooks and rules. See Table 13 on page 32 .
Administration	Access tasks related to administration, such as managing users and roles, authentication, and licenses. See Table 14 on page 33 .
ATOM	Access the Anuta ATOM application. See " Integration with Anuta ATOM Overview " on page 90.

NOTE: The Paragon Automation suite contains a built-in element management system (EMS) that provides features that are used by one or more Paragon Automation applications. Therefore, in the sections that follow, for the menu items used to access these features, the applicable application is listed as *base component*.

Monitoring Sub-Menu

[Table 9 on page 29](#) displays the menu entries in the Monitoring sub-menu and the applications to which each menu entry is applicable.

Table 9: Monitoring Sub-Menu Entries

Sub-menu Entry and Applicable Applications	Description
Network Health (Insights)	See "About the Network Health Page" on page 833.
Alarms and Alerts > Alarms (Insights)	See "About the Alarms Page" on page 847.
Alarms and Alerts > Alerts (Insights)	See "About the Alerts Page" on page 850.
Graphs > Charts (Insights)	See "Manage the Charts Page" on page 404.
Graphs > Grafana	See "Grafana Overview" on page 423
Jobs (base component)	See "About the Jobs Page" on page 854.
Logs	
Action Engine	See "About the Workflows Monitor Page" on page 900.

Reports Sub-Menu

[Table 10 on page 29](#) displays the menu entries in the Reports sub-menu and the applications to which each menu entry is applicable.

Table 10: Reports Sub-Menu Entries

Sub-menu Entry and Applicable Applications	Description
Reports > Health Reports (Insights)	See "About the Health Reports Page" on page 911.
Reports > Network > Integrity Check (Pathfinder)	See "View Network Reports" on page 915.
Reports > Network > LSP Discrepancy (Pathfinder)	

Table 10: Reports Sub-Menu Entries (Continued)

Sub-menu Entry and Applicable Applications	Description
Reports > Maintenance > Link Oversubscription (Pathfinder)	See " View Maintenance Reports " on page 919.
Reports > Maintenance > Link Utilization Changes (Pathfinder)	
Reports > Maintenance > LSP Path Changes (Pathfinder)	
Reports > Maintenance > Maintenance Simulation (Pathfinder)	
Reports > Maintenance > Path Delay (Pathfinder)	
Reports > Maintenance > Peak Interface Utilization (Pathfinder)	
Reports > Maintenance > Peak Link Utilization (Pathfinder)	
Reports > Maintenance > Peak Simulation Summary (Pathfinder)	
Reports > Maintenance > Peak Tunnel Failure (Pathfinder)	
Reports > Inventory (Pathfinder)	
Reports > Demand (Pathfinder)	" View Demand Reports " on page 948

Network Sub-Menu

[Table 11 on page 31](#) displays the menu entries in the Network sub-menu and the applications to which each menu entry is applicable.

Table 11: Network Sub-Menu Entries

Sub-menu Entry and Applicable Applications	Description
Topology (Pathfinder)	See "About the Topology Page" on page 677.
Traffic Engineering > Bandwidth Calendar (Pathfinder)	See "About the Bandwidth Calendar Page" on page 819.
Traffic Engineering > Path Optimization (Pathfinder)	See "About the Path Optimization Page" on page 821.
Change Control Management (Pathfinder)	See "About the Change Control Management Page" on page 829.

Planning Sub-Menu

[Table 12 on page 31](#) displays the menu entries in the Planning sub-menu and the applications to which each menu entry is applicable.

Table 12: Planning Sub-Menu Entries

Sub-menu Entry and Applicable Applications	Description
Paragon Planner	See "Access Paragon Planner Desktop Application" on page 48. See "Access Paragon Planner Web Application" on page 49.

Configuration Sub-Menu

[Table 13 on page 32](#) displays the menu entries in the Configuration sub-menu and the applications to which each menu entry is applicable.

Table 13: Configuration Sub-Menu Entries

Sub-menu Entry and Applicable Applications	Description
Devices (base component)	See "About the Devices Page" on page 138.
Device Groups (Insights)	See "About the Device Groups Page" on page 175.
Network Groups (Insights)	See "About the Network Groups Page" on page 264.
Templates > Config Templates (base component)	See "About the Configuration Templates Page" on page 281.
Templates > Device Templates (base component)	See "About the Device Templates Page" on page 299.
Playbooks (Insights)	See "Manage Playbook Instances" on page 312.
Rules (Insights)	See "About the Rules Page" on page 340.
Resources (Insights)	"About the Resources Page" on page 372
Data Ingest > Settings (Insights)	See "About the Ingest Settings Page" on page 454.
Data Ingest > Diagnostics (Insights)	See "About the Diagnostics Page" on page 561.
Insights Settings (Insights)	See "About the Insights Settings Page" on page 613.
Summarization Profiles > Raw Data (Insights)	See "About the Raw Data Summarization Profiles Page" on page 649.
Summarization Profiles > Data Roll Up (Insights)	See "About the Data Roll Up Summarization Profiles Page" on page 650.
Device Images (base component)	See "About the Images Page" on page 190.

Table 13: Configuration Sub-Menu Entries (Continued)

Sub-menu Entry and Applicable Applications	Description
Network Settings > Topology Filter (Pathfinder)	See "About the Topology Filter Page" on page 270.
Network Settings > Admin Group (Pathfinder)	See "Assign Names to Admin Group Bits" on page 197.
Network Settings > Pathfinder (Pathfinder)	See "Modify Pathfinder Settings From the Pathfinder CLI" on page 198 and "Modify Pathfinder Settings From the GUI" on page 206.
Action Engine (Insights)	See "About the Workflows Page" on page 902.

Administration Sub-Menu

[Table 14 on page 33](#) displays the menu entries in the Administration sub-menu and the applications to which each menu entry is applicable.

Table 14: Administration Sub-Menu Entries

Sub-menu Entry and Applicable Applications	Description
User Management > User (base component)	See "About the Users Page" on page 58.
User Management > Role (base component)	See "About the Roles Page" on page 67.
User Management > User Groups (base component)	See "About the User Groups Page" on page 75.
Authentication > Portal Settings (base component)	See "Configure Portal Settings" on page 55.
Authentication > SMTP Settings (base component)	See "Configure SMTP Settings" on page 50.
Authentication > LDAP Settings (base component)	See "Configure LDAP Settings" on page 53.

Table 14: Administration Sub-Menu Entries (Continued)

Sub-menu Entry and Applicable Applications	Description
Authentication > Email Templates (base component)	See "About the E-mail Templates Page" on page 954.
Authentication > Identity Providers (base component)	See "About the Identity Providers Page" on page 81.
Audit Logs (base component)	See "About the Audit Logs Page" on page 959.
External EMS (base component)	See "About the External EMS Page" on page 963.
Task Scheduler (Pathfinder)	See "About the Task Scheduler Page" on page 969.
Security (Insights)	See "About the Security Settings Page" on page 1001.
License Management (base component)	See "About the License Management Page" on page 1006.

Access the Paragon Automation GUI

IN THIS CHAPTER

- [Single Sign-On Overview | 35](#)
- [Access the Paragon Automation GUI | 36](#)
- [About the Dashboard Page | 38](#)
- [Change Your Password | 45](#)
- [Reset Your Password | 45](#)
- [Reset User Password Through System Console | 46](#)

Single Sign-On Overview

Paragon Automation allows single sign-on of users to allow users to access Paragon Automation resources by using credentials of a third-party account such as Anuta or Google.

NOTE: Paragon Automation supports single sign-on only for Anuta and Google accounts.

When you log in to Paragon Automation by using the credentials of an identity provider account, the authentication is provided by the identity provider and authorization by Paragon Automation.

To allow single sign-on in Paragon Automation:

1. Add the identity provider to Paragon Automation; see ["Add Identity Providers" on page 84](#).
2. If default roles are not assigned for users logging in by using the identity provider account credentials while adding the identity provider, add the default roles; see ["Edit and Delete Identity Providers" on page 87](#).

You can also define individual users who can log in by using single sign-on and assign specific roles to them; see ["Add Users" on page 62](#).

If you have roles defined as part of identity provider and also as an individual user, you are assigned privileges based on the roles assigned to you as an individual user.

After you create an identity provider, Paragon Automation displays **Log in with <identity provider-name>** on the log in page. Users having an account with the identity provider can log in to Paragon Automation by using their credentials of their identity provider account.

RELATED DOCUMENTATION

[About the Identity Providers Page | 81](#)

[Add Anuta ATOM as an Identity Provider | 97](#)

[Integration with Anuta ATOM Overview | 90](#)

Access the Paragon Automation GUI

Before you access the GUI, activate your account.

1. Click the **Set your password** link in the activation mail.

The Set Password page appears.

2. In the Password field, enter your password.

The password should be between 6 to 20 characters and must be a combination of uppercase letters, lowercase letters, numbers, and special characters.

3. In the Confirm Password field, enter the password again for confirmation. .

4. Click **OK**.

An e-mail with the subject Paragon Account Password Changed is sent to you.

Your account is now activated. You can access the Paragon Automation GUI by using the URL and the username present in the activation mail and the password you set.

After an account is created for a user in Paragon Automation, an activation e-mail with a subject Paragon Account Created is sent to the user if SMTP is configured on Paragon Automation. The activation e-mail contains the URL to access the Paragon Automation GUI, the username and the **Set your password** link to set your password. Your user account is activated only after you click the Set your password link and set your password.

If SMTP is not configured, you will be intimated about the URL to access Paragon Automation, your username and password, either verbally or through an e-mail by the system administrator who manages the Paragon Automation installation.

To access the Paragon Automation GUI:

1. Click the URL of the Paragon Automation installation.

The Paragon Automation Login page appears.

2. Enter your username and password.

Username is case insensitive.

On successful authentication, the Dashboard page appears. The navigation menu on the left-hand side of every page allows you to access different objects and perform different tasks easily. The top-level menu items in the navigation menu are listed in [Table 15 on page 37](#).

You need a license to activate the graphical user interface (GUI). Navigate to **Administration > License Management** to add a license. After you successfully add a license for a component (Paragon Insights, Paragon Pathfinder, or Paragon Planner), you can see the related GUI pages. The availability of features in Paragon Insights, Paragon Pathfinder, and Paragon Planner is based on the license you have installed.

The Paragon Automation icon in the top-left corner of the GUI is updated depending on the license that you add. When you log in to the Paragon Automation GUI for the first time, the Paragon Automation icon appears without the names of the three components below it. The GUI displays the name of a component only after you add a license for that component. For example, after you add a Paragon Insights license, the name **Insights** appears below the Paragon Automation icon. After you add a license for Paragon Pathfinder, the names **Pathfinder** and **Planner** are also displayed.

[Table 15 on page 37](#) lists the top level menu items in the Paragon Automation GUI.

Table 15: Paragon Automation GUI Menu

Menu	Description
Favorites	Lists the pages that you marked as favorite. You do not have any listings when you log in to Paragon Automation for the first time.
Dashboard	Displays a variety of status and statistics information related to the network as widgets. These widgets are displayed in a carousel, that you can arrange.
Monitoring	View network health report, alerts and alarms, scheduled jobs, jobs that are initiated and completed, inventory reports, reports about link utilization changes, peak interface utilization and many more reports.
Network	View and manage the network topology, and perform tasks such as configuring tunnels, optimizing paths, and many more tasks.

Table 15: Paragon Automation GUI Menu (Continued)

Menu	Description
Planning	Access the Paragon Planner desktop application to simulate your network.
Configuration	Configure devices, device groups, network groups, configuration templates, device templates, playbooks for analytics, upload device images, and many more objects.
Administration	Perform administration tasks such as configuring security, SMTP settings, LDAP settings, managing users, roles, user groups, licenses, and many more tasks.
ATOM	Access the Anuta ATOM GUI.

RELATED DOCUMENTATION

[Paragon Automation Overview | 3](#)

[Paragon Automation GUI Overview | 18](#)

[Paragon Automation GUI Menu Overview | 27](#)

About the Dashboard Page

IN THIS SECTION

● [Tasks You Can Perform | 39](#)

● [Field Descriptions | 40](#)

To access the dashboard, select **Dashboard**.

Use the Dashboard page to view a variety of status and statistics information related to the network. This information is displayed in a carousel as widgets that you can arrange according to your preference.

[Table 16 on page 39](#) describes the meaning of the severity level colors displayed on the widgets.

Table 16: Meaning of the Severity Level Colors

Color	Definition
Green	Healthy—The overall health of the device, device group, or network group is normal. No problems have been detected.
Yellow	Warning—There might be a problem with the health of the device, device group, or network group. A minor problem has been detected. Further investigation is required.
Red	Risk—The health of the device, device group, or network group is severely compromised. A major problem has been detected.
Gray	No Data—No data is available for the device, device group, or network group.

Tasks You Can Perform

- Show or hide the carousel—Click the cluster of nine blue dots on the top-center part of the page to display the carousel that contains the available widgets. To hide the carousel, click the up arrow on the carousel.
- View widgets in the carousel—From the list at the top left of the carousel, select the category of widgets that you want to view. The default is **All Widgets**.

The widgets that belong to the selected category are displayed in the carousel.

- Search for a widget in the carousel—Click the search icon (magnifying glass) at the top-left corner of the carousel, enter the search text, and press **Enter**.

The widgets, relevant to your search text, are displayed in the carousel.

- Create your own dashboard—Click the add (+) icon to the right of *My Dashboard* to create your own dashboard.

A new tab appears next to *My Dashboard*.

Then, drag the widgets (that you want to add to your dashboard) from the carousel and drop them in this tab. The dashboard automatically adjusts the placement of the widgets to dynamically fit on your browser window without changing their order. You can press and hold the top portion of the widget and move it to a new location on the tab.

- Update information on specific widgets or all widgets—Click the refresh icon (circular arrow) on the top of a widget to refresh this specific widget or click the refresh icon on the top-right corner of the dashboard to refresh all widgets in the dashboard.

At the bottom of each widget, you can see the date (in MM:DD:YYYY format) and time (in HH:MM:SS AM/PM 12-hour format) at which the widget was last updated.

- Rename a dashboard—Double-click on the title bar of the dashboard, specify a name, and press **Enter**.
- Delete a dashboard or widget—Click the delete icon (X) in the title bar of the dashboard to remove the dashboard from the Dashboard page or click X on the top of a widget to remove the widget from the Dashboard page, and confirm the delete operation. The dashboard or widget, and its custom settings (if any) are deleted.

Field Descriptions

Table 17 on page 40 describes the dashboard widgets.

Table 17: Widgets on the Dashboard

Widget	Description
<i>Alerts</i>	
Alert Severity	<p>Displays a donut chart that is segmented by severity level of the alert (Normal, Major, or Minor).</p> <p>The integer value at the center of the donut chart indicates the total number of alerts. Hover over a segment to view the number (and percentage share) of alerts that belong to a severity level.</p>
Alert Summary	<p>Displays a graph with timeline view that indicates the number of alerts that belong to each severity level (Normal, Major, or Minor).</p> <p>Hover your cursor over any time point in the chart to see the number of alerts that belong to each severity level for that time.</p>
Alert Status	<p>Displays a donut chart that is segmented by alert status (Open, Closed, or Expired).</p> <p>The integer value at the center of the donut chart indicates the total number of alerts. Hover over a segment to view the number (and percentage share) of alerts that are open, closed, or expired.</p>

Table 17: Widgets on the Dashboard (*Continued*)

Widget	Description
Alarms By Type	<p>Displays a bar graph that indicates the number of major and minor alarms, per alarm type.</p> <p>From the View list at the top-right corner of the widget, select the period for which you want to view the alarms. You can choose to view alarms for the past hour, past day, past week, and past month.</p> <p>To see more details of the alarms (such as the severity level and time at which the alarm was raised), click the <i>More Details</i> link at the bottom-right corner of the widget. You are taken to the Alarms page (Monitoring > Alarms).</p>
Historical Alarms	<p>Displays a graph with timeline view that indicates the number of major and minor alarms for the selected period.</p> <p>From the View list at the top-right corner of the widget, select the period for which you want to view the alarms. You can choose to view alarms for the past hour, past day, past week, past month, and past year.</p> <p>To see more details of the alarms (such as the severity level and time at which the alarm was raised), click the <i>More Details</i> link at the bottom-right corner of the widget. You are taken to the Alarms page (Monitoring > Alarms).</p>
<i>Devices</i>	
Device Count By Vendor	<p>Displays a donut chart that is segmented by vendor.</p> <p>The integer value at the center of the donut chart indicates the total number of devices in the network. Hover over a segment to view the number (and percentage share) of devices belonging to a particular vendor.</p>
Top 10 Traffic - Devices	Displays a bar graph that indicates the top 10 devices (in descending order) with the highest measured traffic values, in terms of the bandwidth usage.
Top 10 Traffic - Interfaces	Displays a bar graph that indicates the top 10 interfaces (in descending order) with the highest measured traffic values, in terms of the bandwidth usage.
Top 10 Delay - Interfaces	Displays a bar graph that indicates the top 10 interfaces (in descending order) with the highest measured delay.

Table 17: Widgets on the Dashboard *(Continued)*

Widget	Description
Devices	<p>Displays a table that lists all the devices in the network and their details (status, host IP address, and the number of device groups that each device belong to).</p> <p>If you click a device name, you are taken to the Devices page (Configuration > Devices), where you can perform various actions on the device. If you click a device group value, you are taken to the Device Group Configuration page (Configuration > Device Group) where you can perform various actions on the device group.</p>
<i>Status</i>	
Device Status	<p>Displays a donut chart that is segmented by the health status of devices (Healthy, Warning, Risk, or No Data).</p> <p>The integer value at the center of the donut chart indicates the total number of devices in the network. Hover over a segment to view the number (and percentage share) of devices with a particular health status.</p>
Device Group Health	<p>Displays a donut chart that is segmented by device groups.</p> <p>The integer value at the center of the donut chart indicates the total number of devices in device groups. Each segment on the donut chart represents a device group, while the color of each segment represents the health status (Healthy, Warning, Risk, or No Data) of the device group. Hover over a segment to view the number (and percentage share) of devices in a device group with a particular health status.</p>
Network Health	<p>Displays a donut chart that is segmented by network groups.</p> <p>The integer value at the center of the donut chart indicates the total number of networks in the network group. Each segment on the donut chart represents a network group, while the color of each segment represents the health status (Healthy, Warning, Risk, or No Data) of the network group. Hover over a segment to view the number (and percentage share) of devices with a particular health status.</p>

Table 17: Widgets on the Dashboard (Continued)

Widget	Description
Device Groups	<p>Displays separate panels for each device group.</p> <p>Each panel displays:</p> <ul style="list-style-type: none"> • The number of devices in a device group. Click the down arrow beside the device count to see the list of devices. • The number of playbooks to which the device group is applied. Click the down arrow beside the playbook count to see the list of playbooks. • A donut chart that is segmented by the health status for the device group. The integer value at the center of the donut chart indicates the total number of devices in the device group. Hover over a segment to view the number (and percentage share) of devices that belong to a particular health status. <p>To add a device group to the widget, click the <i>Add New Device Group</i> link that appears on the empty panel in the widget. You are taken to the Device Group Configuration page (Configuration > Device Group), where you can add a device group.</p>
Network Groups	<p>Displays a table with the health status for each network group.</p> <p>Click a network name to view the network properties (such as the number of playbooks and instances for the network) on the Network Configuration page (Configuration > Network Groups).</p>
<i>LSPs</i>	
Top LSP Sources	<p>Displays a table that lists the top 10 routers that have LSPs originating there, and the number of originating LSPs, based on LSP count.</p>
Top LSP Destinations	<p>Displays a table that lists the top 10 routers that have LSPs terminating there, and the number of terminating LSPs, based on LSP count.</p>
LSP Summary	<p>Displays a bar graph that indicates the number of primary, standby, and secondary LSPs that are Up (or Active) and Down.</p>

Table 17: Widgets on the Dashboard *(Continued)*

Widget	Description
LSP Hop Count	Displays a bar graph that indicates the number of LSPs by hop count, per LSP control type (PCE-initiated, Delegated, and Device-controlled).
Top 10 Traffic - LSPs	Displays a bar graph that indicates the top 10 LSPs (in descending order) with the highest measured traffic values, in terms of the bandwidth usage.

TSDB (Time Series Database)

TSDB Buffer Bytes	Displays a graph with timeline view that indicates the buffered bytes per TSDB node. Hover your cursor over any time point in the chart to see the buffered bytes per node for that time.
TSDB Buffer Length	Displays a graph with timeline view that indicates the buffer length per TSDB node. Hover your cursor over any time point in the chart to see the buffer length per node for that time.
TSDB Read Errors Last 5min	Displays a bar graph that indicates the number of TSDB read errors, per node, reported in the past 5 minutes.
TSDB Write Errors Last 5min	Displays a bar graph that indicates the number of TSDB write errors, per node, reported in the past 5 minutes.
Latest TSDB Buffer Length	Displays a bar graph that indicates the TSDB buffer length, per node.

RELATED DOCUMENTATION

[Paragon Automation GUI Overview | 18](#)

[Paragon Automation GUI Menu Overview | 27](#)

Change Your Password

Your password expires 180 days after it is assigned to you. After your password expires, you will not be able to log in to Paragon Automation. So, you must change your password before it expires.

To change your password:

1. Log in to the Paragon Automation GUI.
2. On the top-right corner, on the banner of the Paragon Automation GUI, click the user icon and select **Change Password**.

The Change Password page appears.

3. In the Current Password field, enter your current password.
4. In the New Password field, enter a new password or select the **Use a Securely Generated Password** option that appears when you click in the field.

The password should be between 6 to 20 characters and must be a combination of uppercase letters, lowercase letters, numbers, and special characters.

A password strength indicator indicates the strength of your new password. A long green line indicates a strong password where as a short red line indicates a weak password. We recommend that you provide a strong password.

5. In the Confirm Password field, enter the new password again for confirmation.

If you selected the Use a Securely Generated Password option for entering new password, the Confirm Password field is auto-populated.

6. Click **OK**.

A message indicating that the password is changed appears. Use your new password to log in to Paragon Automation the next time.

RELATED DOCUMENTATION

[Paragon Automation GUI Overview | 18](#)

[Paragon Automation GUI Menu Overview | 27](#)

Reset Your Password

Sometimes you may forget your password. In such a case, you can use the Forgot Password option to reset your old password and obtain a new password.

NOTE: You cannot reset passwords for the predefined user(admin). You must contact your system administrator to obtain or reset passwords for the predefined user.

Users with sp-admin, sp-operator, and other custom roles must change the account passwords using the following procedure every 180 days.

To reset your password:

1. On the Login page of Paragon Automation, in the Username field, enter your username in the user@domain.com format.
2. Click the **Forgot Password** link that appears below the Log In button.
A message appears stating that a mail is sent with a link to reset your password to your e-mail id.
3. Open the mail and click the **Reset Your Password** link.
The Reset Password page appears.
4. In the **New Password** field, enter a new password.
The password should be between 6 to 20 characters and must be a combination of uppercase letters, lowercase letters, numbers, and special characters.

A password strength indicator indicates the strength of your new password. A long green line indicates a strong password where as a short red line indicates a weak password. We recommend that you provide a strong password.

5. In the **Confirm Password** field, enter the new password again for confirmation.
6. Click **OK** to reset your password.
A message indicating that your password is reset appears.

You can now log in to Paragon Automation with your new password.

RELATED DOCUMENTATION

[Paragon Automation GUI Overview | 18](#)

[Paragon Automation GUI Menu Overview | 27](#)

Reset User Password Through System Console

In Paragon Automation, the default username and password are set as *admin* and *Admin123!*, respectively.

The admin user does not require an e-mail address to access the Paragon Automation GUI. If the initial password set by an admin user is lost, it can be recovered by a system administrator who has access to the physical server or virtual machine where Paragon Automation is installed.

To recover the admin password, the system administrator has to run the following curl command in any shell in one of the nodes in the Kubernetes cluster.

NOTE: The system administrator can use the command to reset the password of any user.

```
curl -k --request POST 'https://{{server-ip}}:{{port}}/iam/reset-password' --header 'x-service-token: '$
(kubectl get secret -n {{namespace}} $(kubectl get sa -n {{namespace}} iam -o jsonpath='{.secrets[0].name}')
-o jsonpath='{.data.token}' | base64 --decode)'' --header 'x-service-scope: {}' --header 'Content-Type:
application/json' --data '{
  "user_name" : "{{username}}",
  "new_password" : "{{password}}"'
}'
```

Where,

- server-ip denotes the virtual IP address of the Ingress Controller that you configured during installation.
- Port is 443
- namespace is **common**.
- username is the username of the user.
- password is the new password for the user that you enter.

The password should be between 6 to 20 characters and must be a combination of uppercase letters, lowercase letters, numbers, and special characters.

When the system administrator runs this command, the user's password is reset to the password that the system administrator provided and an e-mail is sent to the user with the new password, if SMTP is configured. If SMTP is not configured, the system administrator must manually inform the user about the new password.

RELATED DOCUMENTATION

[Reset Your Password | 45](#)

[Change Your Password | 45](#)

Access the Paragon Planner

IN THIS CHAPTER

- Access Paragon Planner Desktop Application | 48
- Access Paragon Planner Web Application | 49

Access Paragon Planner Desktop Application

The Paragon Planner desktop application helps you simulate the effect of changes in a network without affecting the actual network. You can access the Paragon Planner desktop application by using the Java Network Launch Protocol (JNLP) file, which you can download from the Paragon Planner web UI. JNLP files are used to launch applications hosted on a web server or a remote computer.

NOTE: To run the JNLP file, you must have Java Runtime Environment installed on your computer.

To download and launch the Paragon Planner desktop application:

1. From the Paragon Automation web UI, select **Planning > Paragon Planner Desktop**.
2. In the Paragon Planner Desktop page that is displayed, enter the memory to be allocated and click **OK**. The default value is 512 MB.
3. Click **Save File** when prompted.

A Java Network Launch Protocol (JNLP) file is downloaded to your computer.

4. Double-click the JNLP file to launch Paragon Planner desktop application.
5. Log in using the credentials for Paragon Planner.

Paragon Planner desktop application main window opens. For information about how to use Paragon Planner desktop application, see [Paragon Planner Desktop Application User Guide](#)

RELATED DOCUMENTATION

Key Paragon Planner Features

Main Window Overview

Access Paragon Planner Web Application

Paragon Planner is a network modeling tool that provides offline visualization of network resources and enables detailed architectural planning of the production networks. Paragon Planner allows you to simulate network changes and other traffic scenarios without affecting the production network. Paragon Planner helps you forecast the effect of network changes and assess the network for potential failure scenarios.

NOTE: The Paragon Planner Web Application is a beta feature. To access Paragon Planner Web Application, contact Juniper Networks Technical Assistance Center (JTAC).

Configure SMTP, LDAP, and Portal Settings

IN THIS CHAPTER

- [Configure SMTP Settings | 50](#)
- [LDAP Authentication Overview | 52](#)
- [Configure LDAP Settings | 53](#)
- [Configure Portal Settings | 55](#)

Configure SMTP Settings

You must configure SMTP in Paragon Automation so that the Paragon Automation users can be notified when their account is created, activated, locked, or when password is changed for their account.

To configure SMTP settings:

1. Click **Administration > Authentication > SMTP Settings** in the left navigation bar.
The SMTP Settings page appears.
2. Configure SMTP settings referring to [Table 18 on page 50](#).
3. Click **Save**.
The SMTP settings are saved.

[Table 18 on page 50](#) lists the fields on the SMTP Settings page.

Table 18: SMTP Settings

Field	Description
SMTP Server	
Server Address	Enter the hostname for the SMTP server.

Table 18: SMTP Settings (Continued)

Field	Description
TLS	Click to enable or disable Transport Layer Security (TLS) protocol. If you enable TLS, the e-mail messages are transmitted over an encrypted channel.
Port Number	Enter the port number for the SMTP server. Check with your e-mail service provider for the SMTP port number. By default, the port number is set to 587 when TLS is enabled and to 25 when TLS is disabled. However, you can modify the port number.
SMTP Authentication	
SMTP Authentication	<p>Enable (default) this option if the e-mail server requires authentication before an e-mail can be sent.</p> <p>The Username and Password fields are displayed when you enable this option.</p> <p>Disable this option if you want to configure an unauthenticated e-mail server.</p> <p>The From Name and From Email Address fields are displayed when you disable this option.</p>
From Name	Enter the name to be displayed in the From field in the e-mail sent to a user.
From Email Address	Enter the e-mail address from which Paragon Automation should send e-mails to users. The name can contain only alphanumeric characters and hyphen (-); 35 characters long.
User Name	Configure the username for logging in to the SMTP server.
Password	<p>Enter the password that you want to use for authentication.</p> <p>The password should be between 8 to 20 characters and must be a combination of uppercase letters, lowercase letters, numbers, and special characters.</p>
Confirm Password	Reenter the password for confirmation.
From Name	Enter a name. This name will appear as the sender's name in the e-mails sent to users from Paragon Automation. ;

Table 18: SMTP Settings (*Continued*)

Field	Description
From E-Mail Address	Enter your e-mail address in the user@domain format. This e-mail address will appear as the sender's e-mail address to the e-mail recipient.
Test SMTP Settings	
E-mail Address	Enter your e-mail address to test the SMTP configuration.
Send Test E-mail	Click the Send Test E-mail button. If the settings are correct, you will receive an e-mail from the address configured in the From E-Mail Address field.

RELATED DOCUMENTATION

[Access the Paragon Automation GUI | 36](#)

[About the E-mail Templates Page | 954](#)

LDAP Authentication Overview

LDAP users can log in to Paragon Automation using their LDAP credentials. You can use Active Directory installed on Windows Server 2012 R2 or OpenLDAP version 2.4 for implementing LDAP in Paragon Automation. To facilitate authentication by LDAP, Paragon Automation maps the LDAP user groups to user groups created within Paragon Automation.

A typical workflow of LDAP-based authentication involves the following steps:

1. An LDAP administrator configures LDAP group in an external server and adds users to the LDAP group.
2. The Paragon Automation administrator configures LDAP settings (for example, LDAP server address, SSL certificate, port number to be used to connect with the LDAP server for SSL communication, and so on) in Paragon Automation; see "[Configure LDAP Settings](#)" on page 53.
3. The Paragon Automation administrator adds a user group for LDAP users in Paragon Automation and then:

NOTE: The Paragon Automation administrator may or may not be the same as the LDAP administrator.

- maps the user group to the LDAP user group.

NOTE: The value of the Mapping Provider Group attribute should be the same in both Paragon Automation and LDAP server.

- assigns roles to the user group for authorization.
- assigns users to that user group.

See ["Add User Groups."](#) on page 77

When an LDAP user logs in to Paragon Automation by using their LDAP credentials, Paragon Automation sends a request for authenticating the user to the LDAP server. After the LDAP server successfully authenticates the user, Paragon Automation enforces access control on the user based on the roles that the Paragon Automation administrator previously assigned for the user group.

RELATED DOCUMENTATION

| [User Groups Overview](#) | 74

Configure LDAP Settings

To configure LDAP settings in Paragon Automation:

1. Click **Administration > Authentication > LDAP Settings** in the left navigation menu.

The LDAP Settings page appears.

2. Configure LDAP by referring to [Table 19 on page 54](#).

3. Click **Save**.

LDAP is configured in Paragon Automation.

[Table 19 on page 54](#) describes the attributes in the LDAP configuration settings.

Table 19: LDAP Configuration Settings

Attributes	Description
<i>LDAP Server</i>	
Server Address	<p>Enter the LDAP server URL.</p> <p>For example, ldap.example.net.</p>
SSL	<p>Click to enable or disable SSL for communication between Paragon Automation and the LDAP server.</p>
SSL Certificate	<p>If you enable SSL, enter or browse for the SSL certificate.</p>
Port Number	<p>Enter the port number for connecting with the LDAP server.</p> <p>The default port number if SSL is enabled is 636 and without SSL is 389.</p> <p>Click Test Connection to test the connection with the LDAP server. The server address and port number are configured correctly if you receive a message indicating that the connection with the LDAP server is successful.</p>
<i>LDAP Authentication</i>	
Authentication Method	<p>The authentication method for the LDAP server. It is set to <i>Simple</i> and cannot be edited.</p>
Base Domain Name	<p>Enter the domain name that constitutes the search base for querying the LDAP server.</p> <p>For example: dc=mycompany, dc=net/com.</p>
Bind Domain Name	<p>Enter the user name configured for LDAP authentication.</p> <p>For example: user@mycompany.net.</p>
Bind Password	<p>Enter the password for LDAP authentication.</p> <p>Click the Test Authentication button to test LDAP authentication settings. The LDAP settings are correct if you receive a message indicating that the LDAP settings is successfully authenticated.</p>

Table 19: LDAP Configuration Settings (Continued)

Attributes	Description
<i>User Options (Optional)</i>	
User Attribute	Specify the username attribute used for comparing user entries in the LDAP server. The default is sAMAccountName.
User Filter	Specify the attribute to filter data retrieved from LDAP. The default value is objectClass=person.

RELATED DOCUMENTATION

| [User Groups Overview](#) | 74

Configure Portal Settings

After Paragon Automation is installed or upgraded, a user with the SP administrator role must configure the Paragon Automation portal URL in the Paragon Automation GUI. The portal URL is used as follows:

- Paragon Automation includes the portal URL in the e-mail that is sent to a user when the user's account needs to be activated.
- If you want to allow single sign-on (SSO) of Anuta ATOM users into Paragon Automation, you must use the portal URL when you register Paragon Automation with Keycloak.

NOTE: Whenever the URL to access Paragon Automation is changed, you must update the Paragon Automation portal URL in the Portal Settings. This ensures that:

- The activation e-mails sent to new users contain the correct link to access the Paragon Automation GUI.
- Single sign-on users are able to access the Paragon Automation GUI.

To configure portal settings:

1. Click **Administration > Authentication > Portal Settings** in the left navigation pane.

The Portal Settings page appears.

2. In the **Portal URL** text field, modify the URL present for accessing the Paragon Automation GUI. The URL must be in one of the following formats:

- `https://hostname.domain-name`, where *hostname* is hostname on which Paragon Automation is installed, and *domain-name* is the domain assigned to the server; for example, pa-server.com.
- `https://server-ipv4-address`, where *server-ipv4-address* is the IPv4 address of the server on which Paragon Automation is installed.

By default, an accessible Web URL or the IPv4 address of the load balancer server is set as the portal URL.

3. Click **Save**.

A message indicating that the Paragon Automation portal URL is updated successfully is displayed.

What's Next

After you update the portal URL in the Paragon Automation GUI, you must inform existing users about the new URL to access Paragon Automation, through e-mail or some other means (for example, Slack).

RELATED DOCUMENTATION

[Integration with Anuta ATOM Overview | 90](#)

[Access the Paragon Automation GUI | 36](#)

Manage Users

IN THIS CHAPTER

- [Users Overview | 57](#)
- [About the Users Page | 58](#)
- [Add Users | 62](#)
- [Edit and Delete Users | 64](#)

Users Overview

Paragon Automation provides a default user, admin, with the sp-admin role. The admin (default user) has full access to Paragon Automation GUI and API capabilities. The admin can add users, custom roles, and user groups.

Paragon Automation authenticates and allows users role-based access control (RBAC) to its resources and actions; that is, you can only access resources and actions that are defined in the roles assigned to you. For example, if you are assigned the sp-operator role, you can only view details of objects (for example, devices, alarms, alerts, device templates, and configuration templates). You do not have the permission to create, add, modify, or delete the objects.

Access controls help you to assign different access privileges to different users.

RELATED DOCUMENTATION

[Roles Overview | 66](#)

[User Groups Overview | 74](#)

About the Users Page

IN THIS SECTION

- [Tasks You Can Perform | 58](#)
- [Field Descriptions | 59](#)

To access this page, click **Administration > User Management > Users**.

Tasks You Can Perform

You can perform the following tasks from this page:

- View details of a user

To view details of a specific user, select the user and click **More > Detail**. Alternatively, hover over the user name and click the **Details** icon that appears.

The Details for *<username>* pane appears on the right side of the page displaying basic information, such as the roles assigned to the user, whether the user is enabled or disabled, provider type of the user and the status of the user.

- Add a User; see ["Add Users" on page 62](#)
- Edit and delete a user; see ["Edit and Delete Users" on page 64](#)
- Resend the activation link.

To resend the activation link, click **More > Resend Activation Link**. The activation mail with the link to activate the user account is resent to the user.

You must resend the activation link if the user does not activate their account within 24 hours.

- Filter Entries—Filter the table entries by adding new filtering criteria.

Hover over the **Filter** (funnel) icon and select **Add Filter**. On the **Add Criteria** page, select the filtering criteria from the **Field** and **Condition** list, and enter the text to be compared in the **Value** field. Click **Add**.

The filtered table entries are listed and the filter criteria name is displayed above the table column names.

You can click the cross (X) icon (next to the filter name) to remove the filtering criteria.

NOTE:

- You can add multiple filtering criteria. Once you add the multiple filtering criteria, you can select the **And** condition to display the entries matching all the filtering criteria or select the **Or** condition to display the entries matching any one of the filtering criteria.
- Quick filter: Save the filtering criteria as quick filters. Once you have added all the filtering criteria, you can save a particular criteria or multiple criteria for future use by clicking **Save**.

On the **Save Filter** window, 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 a specific column in the table.
Hover over the **More Options** (vertical ellipsis) > **Show/Hide Columns** and select the *Column-Name* check box of the columns you want to display in the table.
- Reset Preference—Reset the displayed columns to the default set of columns for each tab in the table.
Hover over the **More Options** (vertical ellipsis) and select **Reset Preference**.
- Sort Entries—Click the column name to highlight the up and down arrows next to the column name. Sort the table entries in ascending or descending order of that column by clicking the up or down arrow respectively.

Field Descriptions

Table 20 on page 59 displays the fields on the Users page.

Table 20: Fields on the Users Page

Field	Description
User Name	The name of the user.

Table 20: Fields on the Users Page (Continued)

Field	Description
Roles	<p>The roles assigned to the user.</p> <p>By default, this column lists only one role assigned to the user. When a user is assigned more than one role, a + <i><integer></i> icon (for example: +2) appears to the right of the role. The integer indicates the number of additional roles assigned to the user. Click on the integer to view the additional roles.</p>
State	<p>Indicates whether the user can log in to Paragon Automation (enabled) or cannot log in to Paragon Automation (disabled).</p>
Provider Type	<p>Indicates who is the provider for authentication and authorization.</p> <ul style="list-style-type: none"> • Local—The authentication and authorization is done within Paragon Automation or by using an LDAP server. • OpenID Connect—A third-party, such as Google or Anuta ATOM, authenticates and authorizes the user.

Table 20: Fields on the Users Page (Continued)

Field	Description
Status	<p data-bbox="544 348 885 373">Indicates a user's account status.</p> <ul data-bbox="544 411 1422 1287" style="list-style-type: none"> <li data-bbox="544 411 1422 474">• Activation_pending: The user account is not activated and therefore the user cannot log in to Paragon Automation. The user must click the activate mail that they received after their account was created and change their password to activate their account. <li data-bbox="544 606 1422 632">• Active: The user account is active. The user can log in to Paragon Automation.: <li data-bbox="544 674 1422 831">• Invite_Expired The invite mail that is sent by Paragon Automation, after the user account was created has expired. The activate mail is valid only for 24 hrs. To activate the user account, the user must contact the administrator to resend a new activation mail. <li data-bbox="544 873 1422 989">• InActive: The user has not logged into Paragon Automation for 180 days. To reactivate the user account, the user must contact the administrator to resend the activation mail. <li data-bbox="544 1031 1422 1188">• Locked: The user account is locked. A user account is locked if the user enters an incorrect password for more than 5 times. To unlock, the user must contact the administrator or use the click the Forgot Password link. <li data-bbox="544 1230 1422 1287">• PasswordExpiry: The user password has expired. The password expires 180 days after it is assigned. <p data-bbox="544 1325 1422 1381">If your password has expired, the user must click the Forgot Password link on the Login page or contact their administrator.</p>

RELATED DOCUMENTATION

[Add Roles | 69](#)

[User Groups Overview | 74](#)

[Single Sign-On Overview | 35](#)

Add Users

You can add several types of users from the **Users** page in Paragon Automation.

An administrator or a user with the privilege to add users can add the following types of users to Paragon Automation:

- Local users, where the user is authenticated and authorized by Paragon Automation.
- Lightweight Directory Access Protocol (LDAP) users, where the user is authenticated by the LDAP server, but is authorized by Paragon Automation.
- Third-party users, where the user is authenticated by OpenID Connect or Google Account, and is authorized by Paragon Automation.

NOTE: Before you add a user, you must configure a URL to access the Paragon Automation GUI from the **Portal Settings** page because this URL is sent in the activation e-mail. For more information, see "[Configure Portal Settings](#)" on page 55.

To add a user:

1. Select **Administration > User Management > Users**.

The Users page appears.

2. Click **Add (+)**.

The Add User page appears.

3. Complete the configuration as described in [Table 21 on page 63](#).

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

4. Click **OK**.

A confirmation message is displayed indicating that the user account is added. The user account is listed on the **Users** page.

After the user is added, if SMTP is configured, Paragon Automation sends the user an activation e-mail that contains a link to activate the user's account. If the user does not click the activation link and set a password within 24 hours, the activation link expires and the account is not activated. To activate the account, the user must resend the activation link by clicking **More > Resend Activation Link**.

Table 21: Fields on the Add User Page

Field	Description
Username	<p>Enter one of the following:</p> <ul style="list-style-type: none"> • A valid e-mail address, if SMTP settings are configured for e-mail verification. The user receives an activation e-mail to the specified e-mail address. • A username of your choice if you want to proceed without e-mail verification.
First Name	<p>Enter the first name of the user as a string of alphanumeric characters and some special characters: underscore (_) and hyphen (-) only. The name can include up to 32 characters. A maximum of 32 characters are allowed.</p>
Last Name	<p>Enter the last name of the user as a string of alphanumeric characters and some special characters: underscore (_) and period (.) only. A maximum of 32 characters are allowed. The name can include up to 32 characters.</p>
State	<p>Click the toggle button to disable the user. By default, the user is enabled.</p> <p>NOTE: Only an enabled user can log in to Paragon Automation.</p>
Provider Type	<p>Select the type of authentication service for the user:</p> <ul style="list-style-type: none"> • Local—Paragon Automation or an LDAP server authenticates the user. • OpenID Connect—A third-party such as Anuta ATOM that uses the authentication services of OpenID Connect authenticates the user. <p>You should select this option if you are adding a user who would use the single sign-on feature to log in to Paragon Automation.</p>

Table 21: Fields on the Add User Page (*Continued*)

Field	Description
Password	<p>Enter a password for the user.</p> <p>The password must be between 6 and 20 characters. A combination of uppercase letters, lowercase letters, numbers, and special characters (symbols) is required.</p> <p>NOTE: This field is displayed only if you chose to install Identity Access Management (IAM) without e-mail verification, that is, by setting the <code>skip_email_verification</code> variable to True during installation. This is helpful when you do not want to use the SMTP service for sharing the password with the user; for example, in a lab setup. When the new user logs in to Paragon Automation for the first time with this password, the user is prompted to set a new password immediately. For more information, see "Access the Paragon Automation GUI" on page 36.</p>
Role	<p>Assign one or more roles to the user.</p> <p>To assign roles, select the roles to be assigned in the left box, and then click >. The selected roles are moved to the right box.</p>

RELATED DOCUMENTATION

[Single Sign-On Overview](#) | 35

[Configure SMTP Settings](#) | 50

[About the Identity Providers Page](#) | 81

Edit and Delete Users

IN THIS SECTION

● [Edit Users](#) | 65

● [Delete Users](#) | 65

Edit Users

To edit or delete a user account, you should be an administrator or a user with the permissions to edit and delete users.

To modify the information about a user:

1. Click **Administration > User Management > User**.

The Users page appears.

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

The Edit User page appears.

3. Modify the parameters by following the guidelines provided in [Table 21 on page 63](#).

4. Click **OK** to save the changes.

If you click OK, a confirmation message indicating that the user account is modified appears and the updated information about the user appears on the Users page.

Delete Users

To delete a user account from Paragon Automation:

1. Click **Administration > User Management > User**.

The Users page appears.

2. Select the user that you want to delete and click the **Delete** (trashcan) icon.

An message appears, asking you to confirm the delete operation.

3. Click **Yes** to delete the users.

A confirmation message indicating that the selected user account is deleted from Paragon Automation appears and the user account is removed from the Users page.

RELATED DOCUMENTATION

[Users Overview | 57](#)

[User Groups Overview | 74](#)

Manage Roles

IN THIS CHAPTER

- [Roles Overview | 66](#)
- [About the Roles Page | 67](#)
- [Add Roles | 69](#)
- [Edit, Clone, and Delete Roles | 71](#)

Roles Overview

IN THIS SECTION

- [Types of Roles | 66](#)
- [Access Privileges | 67](#)

A role is a function assigned to a user that defines the tasks that the user can perform within Paragon Automation. Each user can be assigned one or more roles depending on the tasks that the user is expected to perform. User roles enable you to classify users based on the privileges assigned to them to perform tasks.

Types of Roles

Paragon automation provides the following two types of roles:

- **Predefined roles**—System-defined roles with a set of predefined access privileges. Predefined roles are created while installing Paragon Automation. The predefined roles available are sp-admin and sp-operator.

NOTE: The other predefined roles available are tenant-admin, tenant-operator, opco-admin, and opco-operator. However, these roles are not supported in Paragon Automation.

- Custom roles—User-defined roles with a set of access privileges. Custom roles can be created by the sp-admin or a user with the privilege to create users.

Access Privileges

User roles define the access privileges or permissions and actions to access objects (dashboard, device templates, devices, and so on) in Paragon Automation. For example, a user role can contain permissions to read device configurations and delete alarms and alerts objects.

Paragon Automation provides the following privileges:

- Read
- Create
- Update
- Delete
- Other actions; (for example, stage and deploy for the image objects).

RELATED DOCUMENTATION

[Users Overview | 57](#)

[Add Users | 62](#)

About the Roles Page

IN THIS SECTION

- [Tasks You Can Perform | 68](#)
- [Field Descriptions | 69](#)

To access this page, click **Administration > User Management > Roles**.

Tasks You Can Perform

You can perform the following tasks from this page:

- View details of a role

To view details of a specific role, select the role and click **More > Detail**. Alternatively, hover over the role name and click the **Details** icon that appears.

The Details for *<rolename>* pane appears on the right side of the page displaying basic information, such as the roles scope and a link to the Preview Roles page. The Preview Roles page lists the access privileges assigned to the role.

- Create a custom role; see ["Add Roles" on page 69](#).
- Edit, clone, or delete a role; see ["Edit, Clone, and Delete Roles" on page 71](#).
- Filter Entries—Filter the table entries by adding new filtering criteria.

Hover over the **Filter** (funnel) icon and select **Add Filter**. On the **Add Criteria** page, select the filtering criteria from the **Field** and **Condition** list, and enter the text to be compared in the **Value** field. Click **Add**.

The filtered table entries are listed and the filter criteria name is displayed above the table column names.

You can click the cross (X) icon (next to the filter name) to remove the filtering criteria.

NOTE:

- You can add multiple filtering criteria. Once you add the multiple filtering criteria, you can select the **And** condition to display the entries matching all the filtering criteria or select the **Or** condition to display the entries matching any one of the filtering criteria.
- Quick filter: Save the filtering criteria as quick filters. Once you have added all the filtering criteria, you can save a particular criteria or multiple criteria for future use by clicking **Save**.

On the **Save Filter** window, 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 a specific column in the table.

Hover over the **More Options** (vertical ellipsis) > **Show/Hide Columns** and select the *Column-Name* check box of the columns you want to display in the table.

- **Reset Preference**—Reset the displayed columns to the default set of columns for each tab in the table.

Hover over the **More Options** (vertical ellipsis) and select **Reset Preference**.

- **Sort Entries**—Click the column name to highlight the up and down arrows next to the column name. Sort the table entries in ascending or descending order of that column by clicking the up or down arrow respectively.

Field Descriptions

[Table 22 on page 69](#) describes the fields on the Roles page.

Table 22: Fields on the Roles Page

Field	Description
Role Name	The name of the role.
Role Scope	The scope of the role—Service Provider, Operating Company, and Tenant. NOTE: Operating Company and tenant scopes are not supported in Paragon Automation.
Role Type	The type of role—pre-canned and custom.
Created By	The user who created the role. System indicates that the roles are predefined.

RELATED DOCUMENTATION

[Users Overview | 57](#)

[User Groups Overview | 74](#)

Add Roles

To add a role, you should be an administrator or a user with the privilege to add roles.

To add a role:

1. Select **Administration > User Management > Roles**.

The Roles page appears.

2. Click the **Add (+)** icon to add a new role.

The Add Role page appears.

3. Complete the configuration according to the guidelines provided in [Table 23 on page 70](#).

4. Click **OK**.

A confirmation message indicating that the role is created appears and the role is listed on the Roles page.

[Table 23 on page 70](#) lists the fields on the Add Roles page.

Table 23: Fields on the Add Roles Page

Field	Description
Role Name	Enter a unique name for the role. The name can contain alphanumeric characters, underscore, period, and space; 32-characters maximum.
Description	Enter a description for the role; 255 characters maximum.
Role Scope	Select a scope for the role—Service Provider, Tenant. NOTE: The tenant scope is not supported in Paragon Automation.

Table 23: Fields on the Add Roles Page (Continued)

Field	Description
Access Privileges	<p>Displays the objects in Paragon Automation. You must select the check box against each object and then select the privileges (read, write, update, delete, and other actions) that you want to assign the user for the selected object. You can select one or more access privileges to assign to the user role.</p> <p>NOTE: You must assign at least one access privilege to a role.</p> <p>If you select the first-level objects, the submenu items that belong to the main object and the corresponding access privileges are also selected.</p> <p>The following access privileges can be assigned to a user role:</p> <ul style="list-style-type: none"> • Read— Enables the user to read existing objects. • Create—Enables the user to add new objects. • Update—Enables the user to edit or modify the existing objects. • Delete—Enables the user to delete objects. • Other Actions—Includes actions such as deploy, stage, upload, and simulate.

RELATED DOCUMENTATION

[Add User Groups | 77](#)

[Add Users | 62](#)

Edit, Clone, and Delete Roles

IN THIS SECTION

● [Edit a Role | 72](#)

● [Clone a Role | 72](#)

● [Delete a Role | 73](#)

To edit, clone, and delete roles, you should be an administrator or user with the privilege to edit, clone, and delete roles.

Edit a Role

NOTE: You cannot edit predefined roles.

To edit the parameters configured for a custom role:

1. Select **Administration > User Management > Roles**.

The Roles page appears, displaying the details of the available roles.

2. Select the role that you want to edit and click the **Edit** icon (pencil) to modify the attributes.

The Edit Role page appears. The fields on the Edit Role page are available for editing.

NOTE: You cannot modify the role name and role scope.

3. Modify the role description and privileges as needed.

4. Click **OK** to save the changes.

A message indicating that the role is successfully edited appears and the updated role information is displayed in the Roles table.

Clone a Role

You can clone a role (both custom and predefined) when you want to quickly create a copy of an existing role and modify its access privileges.

To clone a role:

1. Select **Administration > User Management > Roles**.

The Roles page appears, displaying the details of the available roles.

2. Select the role that you want to clone and then click the **Clone** button at the top-right corner of the page.

The Clone Role *Role-Name* page appears.

3. Specify an appropriate name for the cloned role.

The name can contain alphanumeric characters, underscore, period, and space; 32-characters maximum.

4. Click **OK** to save your changes.

A clone of the role is created and listed on the Roles page.

5. Select the new cloned role and click the **Edit** icon (pencil) to modify the parameters.

The Edit Role page appears.

6. Select the objects, and modify the access privileges of the role, as needed.

NOTE: You cannot modify the role name and role scope.

7. Click **OK** to save your changes.

A confirmation message appears, indicating the status of the edit operation.

Delete a Role

You cannot delete a predefined role or if it is assigned to a user.

To delete a role:

1. Select **Administration > User Management > Roles**.

The Roles page appears, displaying the details of the available roles.

2. Select a role that you want to delete and then click the **Delete** (trashcan) icon.

A message appears, asking you to confirm the delete operation.

3. Click **Yes** to delete the selected role.

A confirmation message appears, indicating that the selected role is deleted and the role is no longer listed on the Roles page.

RELATED DOCUMENTATION

[Roles Overview | 66](#)

[Add Users | 62](#)

[Add User Groups | 77](#)

Manage User Groups

IN THIS CHAPTER

- [User Groups Overview | 74](#)
- [About the User Groups Page | 75](#)
- [Add User Groups | 77](#)
- [Edit and Delete User Groups | 79](#)

User Groups Overview

A user group contains a group of users that have the same access privileges assigned to them. A user group helps in grouping users based on their roles and assigning privileges to them easily. You can create user groups that are local to Paragon Automation or create user groups based on user groups present in an LDAP server.

NOTE: Before you create user groups based on user groups in the LDAP server, ensure that you have configured LDAP settings in Paragon Automation; see ["Configure LDAP Settings" on page 53](#) for details.

If you are adding a user group based on a user group in the LDAP server, you must enter the LDAP user group name in the *Mapping Provider Group* field. See ["Add User Groups" on page 77](#) for adding a user group to Paragon Automation.

RELATED DOCUMENTATION

[LDAP Authentication Overview | 52](#)

[Users Overview | 57](#)

About the User Groups Page

IN THIS SECTION

- [Tasks You Can Perform | 75](#)
- [Field Descriptions | 76](#)

To access this page, click **Administration > User Management > User Groups**.

Tasks You Can Perform

You can perform the following tasks from this page:

- View details of a user group

To view details of a user group, click the details icon that appears when you hover over the user group. Alternatively, select the user group and click the **More > Detail** icon. The details is displayed in the Details of *<device group-name>* pane that appears on the right side of the Device Groups page.

- Add a user group; see ["Add User Groups" on page 77](#).
- Edit and delete a user group; see ["Edit and Delete User Groups" on page 79](#).
- Filter Entries—Filter the table entries by adding new filtering criteria.

Hover over the **Filter** (funnel) icon and select **Add Filter**. On the **Add Criteria** page, select the filtering criteria from the **Field** and **Condition** list, and enter the text to be compared in the **Value** field. Click **Add**.

The filtered table entries are listed and the filter criteria name is displayed above the table column names.

You can click the cross (X) icon (next to the filter name) to remove the filtering criteria.

NOTE:

- You can add multiple filtering criteria. Once you add the multiple filtering criteria, you can select the **And** condition to display the entries matching all the filtering criteria or select the **Or** condition to display the entries matching any one of the filtering criteria.

- **Quick filter:** Save the filtering criteria as quick filters. Once you have added all the filtering criteria, you can save a particular criteria or multiple criteria for future use by clicking **Save**.

On the **Save Filter** window, 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 a specific column in the table.

Hover over the **More Options** (vertical ellipsis) > **Show/Hide Columns** and select the *Column-Name* check box of the columns you want to display in the table.

- **Reset Preference**—Reset the displayed columns to the default set of columns for each tab in the table.

Hover over the **More Options** (vertical ellipsis) and select **Reset Preference**.

- **Sort Entries**—Click the column name to highlight the up and down arrows next to the column name. Sort the table entries in ascending or descending order of that column by clicking the up or down arrow respectively.

Field Descriptions

Table 24 on page 76 displays the fields on the User Groups page.

Table 24: Fields on the User Groups Page

Field	Description
Group Name	The name of the user group .
Provider Type	<p>Indicates whether the user group is created within Paragon Automation or in a LDAP server.</p> <ul style="list-style-type: none"> • Local—The user group is created in Paragon Automation. Paragon Automation authenticates and authorizes the user based on roles defined for the user. • LDAP—The user group refers to the user group present in the LDAP server. A user in the user group is authenticated by the LDAP server and authorized by Paragon Automation.

Table 24: Fields on the User Groups Page (Continued)

Field	Description
Users	<p>The users in the user group.</p> <p>By default, this column lists only one user assigned to the user group. When a user group is assigned more than one user, a +<integer> icon (for example: +2) appears to the right of the user name. The integer indicates the number of additional users assigned to the user group. Click on the integer to view the additional users.</p>
Roles	<p>Roles assigned to the user group.</p> <p>By default, this column lists only one role assigned to the user group. When a user is assigned more than one role, a +<integer> icon (for example: +2) appears to the right of the role. The integer indicates the number of additional roles assigned to the user group. Click on the integer to view the additional roles.</p>

RELATED DOCUMENTATION

[LDAP Authentication Overview | 52](#)

[User Groups Overview | 74](#)

[Roles Overview | 66](#)

Add User Groups

To add a user group:

1. Click **Administration > User Management > User Group**.

The User Groups page appears.

2. Click the **Add** (plus) icon.

The Create User Group page appears.

3. Configure values by following the guidelines provided in [Table 25 on page 78](#).

4. Click **OK** to save the changes.

A confirmation message indicating that the user group is successfully created appears and the user group is listed on the User Groups page.

[Table 25 on page 78](#) displays the fields on the Add User Groups page.

Table 25: Fields on the Add User Groups Page

Field	Description
User Group Name	Enter a name for the user group. The name can contain alphanumeric characters, underscore, period, and space; 32-characters maximum.
Provider Type	Select whether the user group should be created in Paragon Automation or referred from an LDAP server. NOTE: If you select LDAP, ensure that you have the LDAP server already configured in Paragon Automation; see "Configure LDAP Settings" on page 53
Mapping Provider Group	If you select LDAP for Provider Type, select the user group or a member from the LDAP server for whom you want to assign the roles.
Roles	Assign roles to the user group by selecting the roles in the left column and clicking > . The selected roles are moved to the right column and assigned to the user group.
Users	If you select Local for Provider Type, select the users and click > to assign roles to the selected users. For users in the local user group, you can provide different roles for different users by selecting the required roles for specific users. For example, user A and user B in a local user group can be assigned the admin and operator role respectively. This assignment is possible if you add user A to the user group and select the sp-admin role from the roles assigned to the user group. Similarly, to assign sp-operator role to user B, add user B to the user group and select the sp-operator role from the roles assigned to the user group . NOTE: For user groups in LDAP server, the selected roles apply to all the users in the group.

RELATED DOCUMENTATION

[LDAP Authentication Overview | 52](#)

[Add Roles | 69](#)

[Single Sign-On Overview | 35](#)

Edit and Delete User Groups

IN THIS SECTION

- [Edit User Groups | 79](#)
- [Delete User Groups | 79](#)

Edit User Groups

To edit a user group, you should be an administrator or a user with the permissions to edit user groups.

NOTE: You cannot edit the name and the provider type of the user group.

To edit the information about a user group:

1. Click **Administration > User Management > User Group**.

The User Groups page appears.

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

The Edit User page appears.

3. Update the users and roles assigned to the user group.

4. Click **OK** to save the changes.

A confirmation message indicating that the user group is modified appears and the updated information about the user group appears on the User Groups page.

Delete User Groups

To delete a user group, you should be an administrator or a user with the permissions to edit and delete user groups.

To delete a user group from Paragon Automation:

1. Click **Administration > User Management > User Group**.

The User groups page appears.

2. Select the user groups that you want to delete and click the **Delete** (trashcan) icon.

An message appears, asking you to confirm the delete operation.

3. Click **Yes** to delete the selected user group.

A confirmation message indicating that the selected user group is deleted from Paragon Automation is displayed and the user group is removed from the User Groups page.

RELATED DOCUMENTATION

[User Groups Overview | 74](#)

[LDAP Authentication Overview | 52](#)

Identity Providers

IN THIS CHAPTER

- [About the Identity Providers Page | 81](#)
- [Add Identity Providers | 84](#)
- [Edit and Delete Identity Providers | 87](#)

About the Identity Providers Page

IN THIS SECTION

- [Tasks You Can Perform | 81](#)
- [Field Descriptions | 83](#)

An identity provider enables you to log in to multiple applications by using the same username and password.

Paragon Automation allows you to use OpenID Connect and Google to provide the authentication services; that is, you can use the username and password of your Anuta ATOM, which uses the OpenID Connect authentication services, or your Google account to log in to Paragon Automation.

To access this page, click **Administration > Authentication > Identity Providers** on the left navigation menu.

Tasks You Can Perform

You can perform the following tasks from this page:

- View details of an identity provider.

To view details of an identity provider, click the details icon that appears when you hover over the identity provider. Alternatively, select the identity provider and click **More > Detail**. The details are displayed in the Details for *<identity-provider>* pane that appears on the right side of the Identity Providers page.

- Add an identity provider; see ["Add Identity Providers" on page 84](#).
- Edit and Delete an identity provider; see ["Edit and Delete Identity Providers" on page 87](#).
- Add custom roles to authorize single sign-on (SSO) users.

You can define custom roles and assign them to users logging into Paragon Automation by using their Anuta ATOM or Google credentials. For information about adding custom roles to SSO users, see ["Edit and Delete Identity Providers" on page 87](#).

- Filter Entries—Filter the table entries by adding new filtering criteria.

Hover over the **Filter** (funnel) icon and select **Add Filter**. On the **Add Criteria** page, select the filtering criteria from the **Field** and **Condition** list, and enter the text to be compared in the **Value** field. Click **Add**.

The filtered table entries are listed and the filter criteria name is displayed above the table column names.

You can click the cross (X) icon (next to the filter name) to remove the filtering criteria.

NOTE:

- You can add multiple filtering criteria. Once you add the multiple filtering criteria, you can select the **And** condition to display the entries matching all the filtering criteria or select the **Or** condition to display the entries matching any one of the filtering criteria.
- Quick filter: Save the filtering criteria as quick filters. Once you have added all the filtering criteria, you can save a particular criteria or multiple criteria for future use by clicking **Save**.

On the **Save Filter** window, 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 a specific column in the table.

Hover over the **More Options** (vertical ellipsis) > **Show/Hide Columns** and select the *Column-Name* check box of the columns you want to display in the table.

- **Reset Preference**—Reset the displayed columns to the default set of columns for each tab in the table.

Hover over the **More Options** (vertical ellipsis) and select **Reset Preference**.

- **Sort Entries**—Click the column name to highlight the up and down arrows next to the column name. Sort the table entries in ascending or descending order of that column by clicking the up or down arrow respectively.

Field Descriptions

Table 26 on page 83 displays the fields on the Identity Providers page.

Table 26: Fields on the Identity Providers Page

Field	Description
Name	The name of the identity provider
Provider Type	The type of identity provider—OpenID Connect, Google.
Issuer	The ID of your identity provider; a URL that uniquely identifies the OIDC identity provider for single sign-on users.
Client ID	The unique ID of Paragon Automation in the authentication server of the identity provider.
Client Secret	The secret generated by the identity provider for Paragon Automation.
Roles	<p>Roles that a user logging into Paragon Automation by using the credentials of an identity provider, can take.</p> <p>This column lists only one role assigned to the identity provider. When an identity provider is assigned more than one role, a +<integer> icon (for example: +2) appears to the right of the role name. The integer indicates the number of additional roles assigned to the identity provider. Click on the <i>integer</i> to view the additional roles.</p>

Table 26: Fields on the Identity Providers Page *(Continued)*

Field	Description
Status	<p>Indicates whether the identity provider is enabled or disabled.</p> <ul style="list-style-type: none"> • Enable: Users can log in to Paragon Automation by using the credentials of the identity provider account. • Disable: Users cannot use the credentials of the identity provider account to log in to Paragon Automation.

RELATED DOCUMENTATION

[Single Sign-On Overview | 35](#)

[Integration with Anuta ATOM Overview | 90](#)

Add Identity Providers

Before you add an identity provider, you must register Paragon Automation with the identity provider. While registering, you must provide the URL where you would be hosting Paragon Automation; see ["Configure Portal Settings" on page 55](#).

To add an identity provider to Paragon Automation, you will need the following information from the identity provider:

- The link to the authentication server of the identity provider (known as issuer).
- Client ID and Client secret.

You can obtain the client ID and client secret details from the identity provider when you register with the identity provider.

Paragon Automation allows you to add OpenID Connect (Anuta ATOM) and Google as identity providers.

To add an identity provider:

1. Select **Administration > Authentication > Identity Providers** on the left navigation menu.

The Identity Providers page appears. If no identity providers are configured, the Identity Providers page has the **Add Identity Provider** button to add an identity provider. If an identity provider is already added, the page lists the details of the identity provider in tabular format.

2. Click **Add Identity Provider** if you are adding an identity provider for the first time or else, click the **Add (+)** icon.

The Add Identity Provider page appears.

3. Enter values as described in [Table 27 on page 85](#).

NOTE: Fields marked with * are mandatory.

4. Click **OK**.

A confirmation message appears indicating that the identity provider is added successfully and the identity provider is listed on the Identity Providers page.

[Table 27 on page 85](#) displays the fields on the Add Identity Providers page.

Table 27: Fields on the Add Identity Providers Page

Field	Description
Type	Select the type of identity provider—OpenID Connect (OIDC) or Google. To add Anuta ATOM as your identity provider, use OpenID Connect.
Name	Enter a name for the identity provider. The name can be a string of alphanumeric characters and some special characters (hyphen and underscore); 32 characters maximum.
Issuer	Enter the URL that uniquely identifies your OIDC identity provider. You can get this URL from the OIDC well-known configuration endpoint. For example: <code>https://sso.server.address.com/.well-known/openid-configuration</code> NOTE: For Anuta ATOM, enter <code>https://<ATOM-PrimaryIP>/auth/realms/system</code> . (Optional) Click Test Connection to verify whether you are able to connect to the issuer. A message indicating that the connection to the identity provider is successful is displayed if Paragon Automation connects with the identity provider.
Status	Click to enable (default) or disable allowing users to log in by using credentials of the identity provider account.

Table 27: Fields on the Add Identity Providers Page (*Continued*)

Field	Description
Client Registration	
Client ID	<p>A unique ID for Paragon Automation in the identity provider. This information is provided by the identity provider.</p> <p>NOTE: For Anuta ATOM, the client ID is derived from ATOM Keycloak UI.</p>
Client Secret	<p>A secret generated for authenticating requests from Paragon Automation. The secret is generated by the identity provider.</p> <p>NOTE: For Anuta ATOM, the client secret is derived from ATOM Keycloak UI.</p>
Authorized Redirect URI	<p>A valid URI pattern, a browser can redirect to after a successful login or logout from Paragon Automation.</p> <p>The value is <code>https://portal ip address/oidc/redirect/callback</code>; where, <i>portal ip address</i> is the IP address to access the Paragon Automation GUI.</p> <p>You cannot edit this field.</p>
Default Role Assignment	
Roles	<p>Assign roles that a user, logging by using the credentials of the identity provider, can take. The left column lists the roles that can be assigned to users.</p> <p>To assign roles, select the roles to be assigned in the left column and then click >. The selected roles are moved to the right column.</p> <p>Users logging in by using the credentials of the identity provider can take up roles present in the right column.</p>

RELATED DOCUMENTATION

[Single Sign-On Overview | 35](#)

[Register Paragon Automation with Keycloak | 93](#)

[Roles Overview | 66](#)

Edit and Delete Identity Providers

IN THIS SECTION

- [Edit Identity Providers | 87](#)
- [Delete Identity Providers | 87](#)

Edit Identity Providers

To edit an identity provider, you should be an administrator or a user with privileges to edit an identity provider. You can edit the following parameters of an identity provider:

- Issuer
- Client ID
- Client Secret
- Role assignment

You can edit the roles assigned for single sign on users.

To edit an identity provider:

1. Select **Administration > Authentication > Identity Providers** on the left navigation menu.

The Identity Providers page appears.

2. Select the identity provider that you want to edit and click the **Edit** (pencil) icon.

The Edit Identity Providers page appears.

3. Edit the fields by referring to [Table 27 on page 85](#).

4. Click **OK**.

A confirmation message appears indicating that the identity provider information is successfully edited and the updated information of the identity provider is listed on the Identity Providers page.

Delete Identity Providers

To delete an identity provider, you should be an administrator or a user with privileges to delete an identity provider.

To delete an identity provider:

1. Select **Administration > Authentication > Identity Providers** on the left navigation menu.

The Identity Providers page appears.

2. Select the identity provider that you want to delete and click the **Delete** (trashcan) icon.

A confirmation message appears.

3. Click **OK**.

A message indicating that the identity provider is successfully deleted appears and the identity provider is no longer listed on the Identity Providers page.

RELATED DOCUMENTATION

[Single Sign-On Overview | 35](#)

[Integration with Anuta ATOM Overview | 90](#)

2

PART

Integration With Anuta ATOM

Integration With Anuta ATOM | 90

Integration With Anuta ATOM

IN THIS CHAPTER

- [Integration with Anuta ATOM Overview | 90](#)
- [Register Paragon Automation with Keycloak | 93](#)
- [Add Anuta ATOM as an Identity Provider | 97](#)
- [Add Anuta ATOM as an External EMS Application | 99](#)
- [Configure Paragon Automation to Route API Calls to Anuta ATOM | 101](#)

Integration with Anuta ATOM Overview

IN THIS SECTION

- [ATOM-Paragon Automation Integration Workflow | 90](#)
- [How ATOM Works with Paragon Automation Components | 91](#)
- [Accessing ATOM from Paragon Automation | 93](#)
- [Benefits | 93](#)

Paragon Automation can integrate with external applications that have EMS and other automation capabilities. Integrating with Anuta Networks ATOM platform is one example of this capability.

ATOM delivers stateful service provisioning and workflow management for stateless services, along with configuration compliance.

ATOM-Paragon Automation Integration Workflow

To integrate ATOM with Paragon Automation:

- Configure the Paragon Automation portal IP address in Paragon Automation.

To configure the Paragon Automation portal IP address, navigate to **Administration > Authentication > Portal Settings**. Enter the IPv4 address for accessing the Paragon Automation GUI in the Authentication Settings page that appears and click **Save**.

- Register Paragon Automation with Keycloak so that ATOM can be added as an external EMS application and identity provider in Paragon Automation; see ["Register Paragon Automation with Keycloak" on page 93](#).
- Add ATOM as an identity provider in Paragon Automation; see ["Add Anuta ATOM as an Identity Provider " on page 97](#).

Adding ATOM as an identity provider allows single sign-on of ATOM users in to Paragon Automation. During the single sign-on by using ATOM credentials, ATOM handles user authentication while Paragon Automation handles authorization.

- Add ATOM as an external EMS application in Paragon Automation; see ["Add Anuta ATOM as an External EMS Application" on page 99](#).

Adding ATOM as an external EMS helps in integrating ATOM with the base component, Paragon Pathfinder, and Paragon Insights components of Paragon Automation.

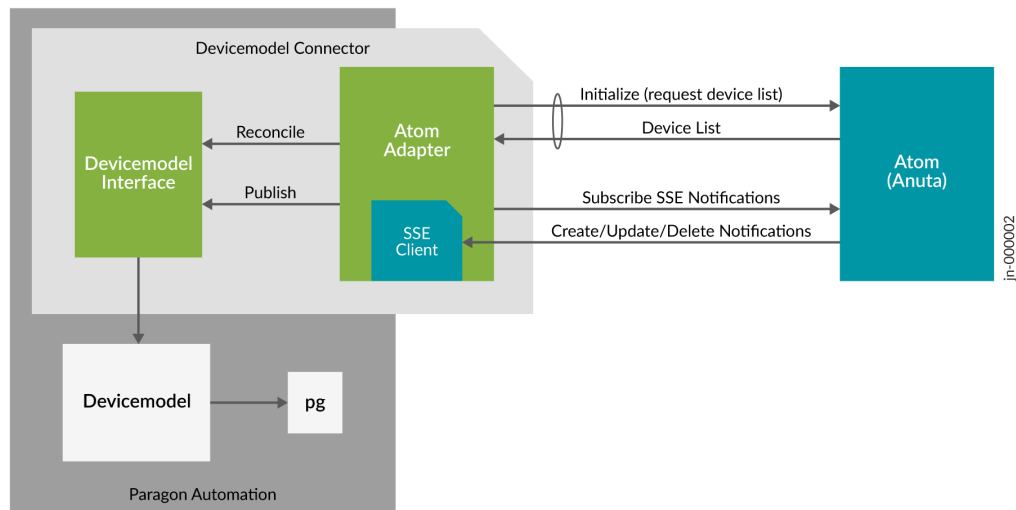
How ATOM Works with Paragon Automation Components

When you add ATOM as an external EMS component in Paragon Automation:

- The device model connector service retrieves the devices discovered by ATOM and synchronizes the device configuration in its database. You can view the devices discovered by both ATOM and Paragon Automation on the Devices page of Paragon Automation.

ATOM sends updates when a device is created, deleted, or modified to Paragon Automation so that the device list in Paragon Automation is synchronized with the network. [Figure 7 on page 92](#) shows the synchronization process for device details.

Figure 7: Device Synchronization



- The Paragon Pathfinder component learns the network topology including all the nodes, links, and label-switched path (LSP) tunnels provisioned on the network. ATOM obtains this information about the LSP tunnels in the network from the Paragon Pathfinder component through REST APIs.

ATOM uses the LSP tunnel information to provision the L3VPN services having different classes of service. If an LSP is not present for configuring an L3VPN with the required class of service, ATOM does not provision the service; instead, it waits for the required LSP to be configured. The service is provisioned after the LSP is configured.

NOTE: ATOM does not configure the LSPs. It merely provides the details (for example, endpoints and bandwidth) for the LSP to be configured. The LSPs must be manually configured by using the Paragon Pathfinder component.

- ATOM integrates with the Paragon Insights component to create and to instantiate playbooks for monitoring the provisioned services.

ATOM monitors the provisioned services for the following key performance indicators (KPIs):

- interface states and errors towards the core edge
- interface states and errors towards the customer edge
- VPN state
- BGP protocol errors

For each service that is provisioned, ATOM creates a playbook instance for monitoring the service and adds the devices on which the service is provisioned to the playbook. Paragon Insights is configured to send notifications through Microsoft Teams, Slack, or e-mail when constraints, such as bandwidth and latency, for the service are not met.

Accessing ATOM from Paragon Automation

When you add ATOM as an identity provider in Paragon Automation, an ATOM user can use the single sign-on feature of Paragon Automation to sign in to Paragon Automation by using their ATOM credentials.

After logging in to Paragon Automation, an ATOM user can log in to ATOM by clicking the **ATOM** link on the left navigation bar. The user is taken to the ATOM portal, where the user can log in by using their ATOM credentials.

Benefits

Integrating Paragon Automation with ATOM provides the following benefits:

- The integration of ATOM with Paragon Automation leverages the traffic engineering (TE) features of Paragon Pathfinder to enable VPN to LSP steering.
- The L3VPN workflows present in ATOM helps in setting up L3 VPNs quickly without errors.
- ATOM configures the L3VPN monitoring playbooks to automate L3VPN assurance.
- The closed-loop automation capabilities of ATOM along with automation capabilities of Paragon Pathfinder and Paragon Insights provide network operators with increased service assurance and operational efficiency.

RELATED DOCUMENTATION

| [Configure Portal Settings](#) | 55

Register Paragon Automation with Keycloak

Anuta ATOM (hereafter ATOM) uses Keycloak for authenticating users. Therefore, to enable ATOM users to use single sign-on to log in to Paragon Automation GUI, you must register Paragon Automation as a client in Keycloak.

NOTE: Ensure that you have the following information before registering Paragon Automation as a client with Keycloak:

- ATOM Keycloak login credentials.
- Paragon Automation portal IPv4 address
- IPv4 address of the ATOM primary node
- Credentials to log in to the ATOM primary node

To register Paragon Automation with Keycloak:

1. Log in to the ATOM Keycloak UI as an admin user by entering the following URL in your browser:
`https://atom-primary-ipv4-address:32443/auth`

NOTE: If using ATOM version 11.1 or later, you can either use `https://atom-primary-ipv4-address:32443/auth` or `https://atom-vip:443/auth`.

2. If the **System** realm is not already added in Keycloak, add it. The System realm is generally pre-created.
3. Add *atom* (generally pre-created) and *paragon-automation* as clients under the *System* realm.

NOTE: For information on how to add realms and clients in Keycloak, refer to the [ATOM API documentation](#).

4. In the left navigation menu, select **System > Clients**.
The Clients page appears.
5. Click the **atom** client.
The atom page appears.
6. On the Settings tab, configure the following parameters:
 - **Valid Redirect URIs:** Enter the following values:
 - `https://paragon-automation-portal-ipv4-address/oidc/redirect/callback`

NOTE: You can also obtain the value for the Valid Redirect Uniform Resource Indicator (URI) from the Authorized Redirect URI field on the Add Identity Providers

page (**Administration > Authentication > Identity Providers > Add Identity Providers**) of the Paragon Automation GUI.

- <https://paragon-automation-portal-ipv4-address>
- <https://paragon-automation-portal-ipv4-address/iam/auth/oidc/callback>
- <https://atom-vip:443/>*

NOTE: If using ATOM version 11.0 or earlier, use <https://atom-primary-ipv4-address:32443/>*

If using ATOM version 11.1 or later, you can either use <https://atom-primary-ipv4-address:32443/>* or <https://atom-vip:443/>*

- Ensure that the following parameters are turned ON:
 - **Enabled**
 - **Direct Access Grants Enabled**
 - **Service Accounts Enabled**
 - **Authorization Enabled**
7. Click **Save**.
 8. Go back to the Clients page.
 9. Click the **paragon-automation** client.
The paragon-automation page appears.
 10. Configure the following parameters:
 - On the Settings tab, ensure that **Enabled** is turned ON.
 - On the Credentials tab:
 - For **Client Authenticator**, select **Client Id and Secret**.
The client secret is displayed in the Secret field. Copy this value because you will need it when you add ATOM as an identity provider in Paragon Automation.
 - For **Valid Redirect URI**, enter the following values:
 - <https://paragon-automation-portal-ipv4-address/oidc/redirect/callback>
 - <https://paragon-automation-portal-ipv4-address>
 - <https://paragon-automation-portal-ipv4-address/iam/auth/oidc/callback>

- `https://atom-vip:443/*`

NOTE: if using ATOM version 11.0 or earlier, use `https://atom-primary-ipv4-address:32443/*`.

If using ATOM version 11.1 or later, you can either use `https://atom-primary-ipv4-address:32443/*` or `https://atom-vip:443/*`.

11. Click **Save** to save the paragon-automation configuration in Keycloak.
12. Log in to the ATOM primary node by using SSH.
13. Open the **oauth2-proxy** deployment file present under the `atom` namespace by executing the following command:

```
kubectl edit deployment -n atom oauth2-proxy
```

The file opens in vi editor.

14. Add the following value to the `- name: OAUTH2_PROXY_EXTRA_JWT_ISSUERS` field:

```
https://atom-vip:443/auth=paragon-automation
```

NOTE: if using ATOM version 11.0 or earlier, use `https://atom-primary-ipv4-address:32443/auth=paragon-automation`.

If using ATOM version 11.1 or later, you can either use `https://atom-primary-ipv4-address:32443/auth=paragon-automation` or `https://atom-vip:443/auth=paragon-automation`.

NOTE: If values are already present in this field, use comma (,) as the separator.

15. Save the file by using the `:wq` command.
16. Exit from the primary node by using the `exit` command.

What's Next

Add ATOM as an identity provider and as an external EMS in Paragon Automation. For more information, see ["Add Anuta ATOM as an Identity Provider" on page 97](#) and ["Add Anuta ATOM as an External EMS Application" on page 99](#).

RELATED DOCUMENTATION

[About the Identity Providers Page | 81](#)

[Single Sign-On Overview | 35](#)

Add Anuta ATOM as an Identity Provider

Before you add Anuta ATOM as an identity provider for single sign-on, you must register Paragon Automation in Keycloak that ATOM uses for authentication. While registering with Keycloak, you must provide the URL to access Paragon Automation portal.

To add ATOM as an identity provider in Paragon Automation, you will need values for Paragon Automation client ID and client secret from Keycloak. The client ID and client secret values are generated in Keycloak when Paragon Automation is registered with Keycloak; see "[Register Paragon Automation with Keycloak](#)" on page 93.

To add ATOM as an identity provider in Paragon Automation:

1. Select **Administration > Authentication > Identity Providers** on the left navigation menu.
The Identity Providers page appears. If no identity provider is already configured, the Identity Providers page has the **Add Identity Provider** button to add an identity provider. If an identity provider is already added, the page lists the configured identity providers.
2. Click **Add Identity Provider** if adding an identity provider for the first time or else, click the **Add (+)** icon.
The Create Identity Provider page appears.
3. Enter values as described in [Table 28 on page 98](#).

NOTE: Fields marked with * are mandatory.

4. Click **OK**.
A confirmation message appears indicating that the identity provider is added successfully and the identity provider is listed on the Identity Providers page.

[Table 28 on page 98](#) displays the fields on the Add Identity Providers page.

Table 28: Fields on the Add Identity Providers Page

Field	Description
Type	Select OpenIDConnect as type of the identity provider.
Name	Enter a name for the identity provider. For example, ATOM The name can be a string of alphanumeric characters and some special characters (hyphen and underscore); 32 characters maximum.
Issuer	Enter the URL in the form https:// <atom-vip:443/auth/realms/system . NOTE: If using ATOM version 11.0 or earlier, use <code>https:// atom-primary-ipv4-address:32443/auth/realms/system</code> . If using ATOM version 11.1 or later, you can either use <code>https:// atom-primary-ipv4-address:32443/auth/realms/system</code> or <code>https:// atom-vip:443/auth/realms/system</code> . (Optional) Click Test Connection to verify whether you are able to connect to the issuer.
Status	If not already enabled, click this toggle button to allow users to log in to Paragon Automation by using ATOM credentials (single sign-on).
Client ID	The unique ID for the Paragon Automation client in the Keycloak. This information is derived from the Keycloak UI. For example: paragon-automation
Client Secret	A secret generated for authenticating requests from Paragon Automation. This information is derived from the Keycloak UI. The secret can be in the format <code>xxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx</code> for the client ID, where x is an alphanumeric character.
Authorized Redirect URI	A valid Uniform Resource Indicator (URI) pattern that a browser can redirect a user to after a successful login or logout from Paragon Automation. The value is <code>https:// portal ip address/oidc/redirect/callback</code> ; where, <i>portal ip address</i> is the IP address to access the Paragon Automation GUI. You cannot edit this field.

Table 28: Fields on the Add Identity Providers Page (*Continued*)

Field	Description
Roles	Assign roles that a user, logging by using single sign-on, can be assigned.

RELATED DOCUMENTATION

[About the Identity Providers Page | 81](#)

Add Anuta ATOM as an External EMS Application

Before you add ATOM as an External EMS application in Paragon Automation, ensure that you have an account created in ATOM and Paragon Automation is registered as a client in Keycloak. To add the ATOM as an external EMS application to Paragon Automation, you need to provide:

- The credentials for accessing ATOM.
This is the username and password set after the ATOM installation, for logging into the ATOM UI. These credentials must be configured in Paragon Automation so that the components in Paragon Automation can interwork with ATOM.
- The client ID and client secret provided by Keycloak when Paragon Automation is registered with Keycloak.

To add ATOM to Paragon Automation as an external EMS application:

1. Click **Administration > External EMS** in the left navigation menu.
The External EMS page appears.
2. Enter values referring to [Table 29 on page 100](#).
3. Click **OK**.

A message indicating whether the External EMS application was successfully added or not is displayed. The External EMS application is listed on the External EMS page if the operation is successful.

[Table 29 on page 100](#) lists the fields to be configured on the Add External EMS page.

Table 29: Fields on the Add External EMS Page

Field	Description
Type	Select ATOM from the list.
Server Address	<p>Enter the IPv4 address to access the ATOM UI.</p> <p>For a non-high availability (non-HA) setup or a high availability (HA) setup, enter the IPv4 address that was configured for the ATOM primary node.</p> <p>For example: 192.168.51.55</p>
Username	Enter admin as the username to log in to ATOM.
Password	<p>Enter the password to log in to ATOM.</p> <p>NOTE: The username and password are configured after ATOM installation.</p>
Parameters	<p>Add parameters, such as client ID and client secret, for connecting with the External EMS application.</p> <p>To add the parameters, click the Add icon (+) provided in the Parameters field and enter the parameter and corresponding value in the text fields that appear. Add more parameters by clicking the Add (+) icon.</p> <p>Some of the parameters that you should configure for adding ATOM as an external EMS application are:</p> <ul style="list-style-type: none"> • client_Id: Client ID of <i>paragon-automation</i> or <i>atom</i> in Keycloak see "Register Paragon Automation with Keycloak" on page 93. • client_secret: Respective client secret of <i>paragon-automation</i> or <i>atom</i> in Keycloak; generally, in the format xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx. • port: Port Paragon Automation uses for connecting with Keycloak. Enter 32443.

NOTE: After adding ATOM as an External EMS to Paragon Automation, ensure that you obtain relevant use case packages required for the Paragon-ATOM integrated Automation solution and install them in ATOM. For information about the packages to be installed for a solution and to access the packages, contact support@anutanetworks.com.

RELATED DOCUMENTATION

| [About the External EMS Page](#) | 963

Configure Paragon Automation to Route API Calls to Anuta ATOM

After Anuta ATOM is added as an external EMS, you can route certain API requests made to Paragon Automation to Anuta ATOM; for example, to access resources in Anuta ATOM through calls made to Paragon Automation APIs. To enable Anuta ATOM to service Paragon Automation API requests, you must configure Anuta ATOM API credentials in Paragon Automation. After the Anuta ATOM API credentials are saved in Paragon Automation, all requests made to `https://paragon/anuta/restconf` are forwarded to `https://anuta/restconf`. You can also configure additional endpoint mappings.

To configure Anuta ATOM API credentials in Paragon Automation, you must configure a Kubernetes secret to provide the credentials. The credentials include `username`, `password`, `client_id`, and `client_secret`.

To configure the Kubernetes secret:

1. Log in to the primary node of Paragon Automation.
2. Create a `*.yml` file with name `anuta-secret` and assign appropriate values for the Anuta ATOM API credentials. The following is a sample of the `anuta-secret.yml` file:

```

...

apiVersion: v1
kind: Secret
metadata:
  name: northstar-anuta-proxy-secret
  namespace: northstar
  labels:
    proxy: anuta-restconf
type: Opaque
stringData:

```

```

username: <username>
password: <password>
client_id: <client_id>
client_secret: <client_secret>
...

```

3. Apply the secret using kubectl:

```
kubectl apply -f anuta-secret.yml
```

4. You can delete the **anuta-secret.yml** file as it is no longer required.

After this configuration, by default, all requests made to <https://paragon/anuta/restconf> is forwarded to <https://anuta/restconf>.

You can also configure additional API endpoint mappings by using Kubernetes ConfigMaps. The following example shows how you can access the Anuta ATOM device resources through Paragon Automation.

To access an Anuta ATOM device resource from Paragon Automation:

1. Log in to the primary node of the Paragon Automation installation.
2. Create a file **anuta-cmap-ctrldev.yml** with a ConfigMap to forward requests matching <https://paragon/restconf/data/controller:devices> to <https://anuta/restconf/data/controller:devices>. The following is a sample of the **anuta-cmap-ctrldev.yml** file.

```

apiVersion: v1
kind: ConfigMap
metadata:
  name: anuta-proxy-restconf-ctrldev
  namespace: northstar
  labels:
    proxy: anuta-restconf
data:
  prefix: /restconf/data/controller:devices
  rewrite: /restconf/data/controller:devices

```

3. For any custom mappings, modify the *prefix* and *rewrite* sections as required, where *prefix* specifies the Paragon Automation URL prefix to match, and *rewrite* specifies the Anuta ATOM URL to where API requests must be forwarded.

4. Apply the ConfigMap using kubectl:

```
kubectl apply -f anuta-cmap-ctrldev.yml
```

After you configure the credentials and endpoints, you can send API requests to a Paragon Automation endpoint. The following example is a request for device information that is fulfilled by Anuta ATOM.

```
curl -k -X GET 'https://<paragon>/restconf/data/controller:devices.json?
fields=device(uuid;device-family-string;device-type;id)' -H "x-iam-token: <token>"

{
  "controller:devices": {
    ...
    "device": [
      {
        ...
        "device-type": "MX960",
        "id": "vmx101",
        "device-family-string": "Juniper MX",
        "uuid": "<uuid>"
      },
      ...
    ]
  }
}
```

By configuring API credentials for Anuta ATOM in Paragon Automation, you can use an API service using Paragon Automation credentials to make requests to both Paragon Automation and Anuta ATOM endpoints.

Delete the ConfigMap if you no longer need the endpoint configuration.

```
kubectl delete -f anuta-cmap-ctrldev.yml
```

RELATED DOCUMENTATION

| [Integration with Anuta ATOM Overview](#) | 90

3

PART

Workflows

[Base Platform](#) | 105

[Paragon Pathfinder](#) | 109

[Paragon Planner](#) | 118

[Paragon Insights](#) | 124

Base Platform

IN THIS CHAPTER

- Onboard and Manage Devices | 105

Onboard and Manage Devices

You can use Paragon Automation to plan, deploy, and monitor your network in both greenfield and brownfield scenarios.

Greenfield deployments involve building a network from the start. In greenfield deployments, you purchase the required number of devices, plan the IP addressing scheme, design network topology, and cable the devices to suit your business requirements.

Brownfield deployments involve provisioning devices that integrate with a legacy network. In brownfield network deployments, you are concerned with scaling up and scaling out devices in your existing network in a fast but efficient manner.

Paragon Automation is a closed-loop automation solution that enables you to automate device provisioning at scale. You can automate device provisioning by using the **Add Devices** or **Discover Devices** option in the Paragon Automation GUI.

Paragon Automation supports selected models in the following device series:

- ACX, MX, and PTX series routers
- EX and QFX series switches
- SRX series firewalls
- Cisco devices
- Nokia devices

See "[Supported Devices](#)" on [page 144](#) for all devices and OS versions supported in Paragon Automation.

TIP: For more information on purchasing a device, go to [How to Buy Juniper Network Devices](#).

After you purchase a device, follow the instructions in the hardware documentation to unbox the device, mount it on a rack, and power on the device. For details about installing a device, see the device's Hardware Guide on the [TechLibrary](#) or the device's [Quick Start Guide](#). Search for the device in the search box provided or click **Routing > View More**, **Switching > View More**, or **Security > View More** and search for the device in the list.

You can onboard devices from Juniper Networks and other vendors using the **Discover Devices** or the **Add Devices (ZTP)** option. Use the **Discover Devices** option to add devices that are already operational in your network (brownfield deployment). Use the **Add Devices (ZTP)** option to add new devices in your network (greenfield or brownfield deployment). During the device onboarding process, Paragon Automation detects the devices using the IP address or hostname that you provide.

Whether you use the **Discover Devices** or the **Add Devices (ZTP)** option to onboard a device, you can mark the management status of devices by using the **Management Status** field. When the **Management Status** field is enabled, devices are managed and when the **Management Status** field is disabled, the devices are unmanaged. Managed devices are those for which Paragon Automation synchronizes the device configuration using NETCONF. Unmanaged devices are those for which Paragon Automation doesn't synchronize the device configuration.

You must ensure that you complete the following device configurations before you add the devices for onboarding:

- Management IP address

NOTE: If your Juniper device has more than one Routing Engine, you must configure an additional management IP address (with the `master-only` statement) that is active only on the primary Routing Engine's interface.

- Configure NETCONF (on port 22) and SSH maximum inbound sessions on Nokia devices
- Configure NETCONF and SSH rate limit on Cisco devices

The workflow for onboarding and managing devices in Paragon Automation is as follows:

1. Onboard the device:

- a. Log in to the Paragon Automation GUI. For more information, see "[Access the Paragon Automation GUI](#)" on page 36.
- b. (Optional) Select **Configuration > Templates > Configuration Templates** to deploy additional configurations such as NTP, Syslog, Aggregated Ethernet device count, SNMP, and so on. See "[Configuration Templates Workflow](#)" on page 280 for more information.
- c. Select **Configuration > Devices**.

The Devices page appears.

- d. Click the **Add (+)** icon.

The Devices page appears.

- e. Select one of the following options to onboard devices in Paragon Automation:

- **Discover Devices:** Use this option if you want to onboard devices that are operational in your network. When you onboard devices using **Discover Devices**, Paragon Automation collects and stores details of the devices in the Paragon Automation database. For more information, see ["Discover Devices" on page 149](#).
- **Add New Devices:** Use this option if you want to onboard new devices by using zero touch provisioning (ZTP). You must specify device details, such as serial number, device model, and the root password. For more information, see ["Add New Devices" on page 153](#).

NOTE: If the devices that you want to onboard (by using ZTP) are not in the same subnet as Paragon Automation, you must install and run DHCP Relay to connect the devices with Paragon Automation. See ["Configure a DHCP Relay for ZTP" on page 134](#) for more information.

Paragon Automation triggers a device discovery job and displays a confirmation message with a link to the job.

- f. Click the job ID link in the confirmation message (or on the Jobs page [**Monitoring > Jobs**]) to open the Job Status page, where you can monitor the status of the scheduled job.
- g. After the job finishes, go to the Devices page and verify that the devices are added or discovered successfully by checking the current status displayed in the Status field.

NOTE:

- For managed devices, the Management Status should be **Up**, indicating that Paragon Automation established a connection with the device. In addition, the Sync Status should be **In Sync**, indicating that the configuration and the inventory data in Paragon Automation and on the device are synchronous.
- For unmanaged devices, the Management Status should be **Unmanaged**, and the Sync Status should be **Unknown**. The Sync Status Unknown indicates that Paragon Automation added the device to its database, but that no NETCONF session was created to synchronize the configuration and the status.

2. (Optional) After you successfully onboard a device, you can perform the following operations to manage a device:

NOTE: For an unmanaged device, you can perform only edit and delete operations.

- Edit the device. See ["Edit Devices" on page 168](#).
- Upgrade the device image. See ["Upgrade the Device Image" on page 156](#).
- View and manage the device configuration. See ["View and Manage Device Configuration" on page 165](#).
- Add configuration templates. See ["Add Configuration Templates" on page 284](#).
- Delete the device. See ["Delete Devices" on page 173](#).

RELATED DOCUMENTATION

| [Add Devices Overview](#) | 135

Paragon Pathfinder

IN THIS CHAPTER

- Acquire and View the Network Topology | 109
- Add LSPs (Tunnels) | 111
- Schedule and Monitor a Maintenance Event | 112
- Reroute LSPs Automatically | 113
- Add and Check Container LSPs | 116

Acquire and View the Network Topology

Paragon Pathfinder acquires the network topology by gathering data from the network by using Border Gateway Protocol Link-State (BGP-LS) and Path Computation Element Protocol (PCEP). You can view the acquired topology on the Topology page (**Network > Topology**) of the Paragon Automation GUI. For more information about viewing and working with the network topology, see ["About the Topology Page" on page 677](#).

Before you configure Paragon Pathfinder to acquire the network topology, you must do the following for the devices to be added to the topology;

- Configure MPLS, RSVP, the interior gateway protocol (IGP) (IS-IS or OSPF) traffic engineering, and NETCONF.
- To view telemetry on the topology, configure:
 - Junos Telemetry Interface (JTI) and Junos Real-time performance monitoring (RPM) service on Juniper devices.
 - OpenConfig or SNMP on Cisco IOS-XR devices.

For details, see the steps to configure the Juniper Networks devices and Cisco devices in the *Workflow to Collect Device Statistics* section in ["Collect Analytics Data Overview" on page 873](#).

To acquire and view the network topology with metrics on the Paragon Automation GUI:

1. Onboard devices to Paragon Automation so that you can start managing the devices; see ["Add Devices"](#) on page 149.
2. For each device that you onboard, edit the fields related to Path Computation Element (PCE) protocol (PCEP), NETCONF, and parameters related to telemetry on the Edit Devices page (**Configuration > Devices > Edit (pencil) icon**) for each device that you added; see ["Edit Devices"](#) on page 168.
3. On the Device Group Configuration page (**Configuration > Device Groups**), verify that the onboarded devices are assigned to the Controller device group.

NOTE: By default, all devices discovered in a topology are added to the Controller device group. However, if you add a device manually, you must add the device to the Controller device group. For more information, see ["Edit a Device Group"](#) on page 183.

4. Set up the device collection task so that you can obtain the configuration on the network devices; see ["Add a Device Collection Task"](#) on page 978.
5. Verify that the status of the device collection task on the Task Scheduler page (**Administration > Task Scheduler**) displays Completed.
6. Configure Paragon Pathfinder to acquire the network topology:
 - a. Configure PCEP on the devices to obtain information about LSPs in the network.
 - b. Enable BGP-LS on the devices to acquire topology.
 - c. (Optional) Configure BGP-LS peers in Paragon Automation if you want to change the BGP-LS peers that you configured while installing Paragon Automation.

For more information, see the *Get Started with Paragon Pathfinder* section of the *Paragon Automation Getting Started Guide* to configure Paragon Pathfinder.

7. Verify that the devices are advertising BGP-LS routes and that Paragon Automation has received the routes.
8. Verify that the network topology is discovered, and that the topology is displayed in the Paragon Automation GUI. On the Topology page:
 - The devices should be displayed with the router icon.
 - Type, IP Address, and Management IP (address) should be displayed for each device on the Node tab of the Network Information table.

RELATED DOCUMENTATION

| [Onboard and Manage Devices](#) | 105

Add LSPs (Tunnels)

Paragon Pathfinder allows you to add label-switched paths (LSPs) by using NETCONF or Path Computation Element Protocol (PCEP). You can add RSVP, segment-routed (SR), or SRv6 LSPs and view the LSPs on the topology map. For more information about how Paragon Pathfinder handles LSPs, see ["Understand How Pathfinder Handles LSPs" on page 715](#) and ["Understand LSP Delegation and Undelegation" on page 815](#).

NOTE:

- LSPs are referred to as tunnels in the Paragon Automation GUI.
- Before you add LSPs, ensure that you have completed the steps to acquire the network topology, as explained in ["Acquire and View Network Topology" on page 109](#).

To add LSPs to your network:

1. Add tunnels (LSPs) on the Topology page (**Network > Topology**) as follows:

- Add a single tunnel; see ["Add a single tunnel" on page 729](#).
- Add diverse tunnels; see ["Add Diverse Tunnels" on page 743](#).
- Add multiple tunnels; see ["Add multiple tunnels" on page 754](#).

After you add the LSPs, the LSPs are listed on the **Tunnels** tab in the Network Information table of the Topology page.

2. 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. (Optional) Use the **Diagnostic** tool present on the top-right side of the Network Information Table to:
 - Check the details of the LSP, as configured on the router at the head-end of the LSP, to confirm that the LSP is configured correctly.
 - Send ping traffic over the LSP to check whether the Node A is able to reach the Node Z.
 - Send traceroute traffic over the LSP to check the intermediate nodes between the Node A and Node Z and any forwarding issues in the LSP.
- d. Check the traffic on and delay in the LSP by selecting the LSP and clicking **View > Tunnel Traffic** and **View > Delay** respectively.

RELATED DOCUMENTATION

[Analyze Root Cause of Network, Device, and Service Issues](#) | 125

Schedule and Monitor a Maintenance Event

Paragon Pathfinder allows you to schedule maintenance events on nodes, links, and facilities (shared risk link groups [SRLGs]). During a maintenance event, Paragon Pathfinder considers the entities (nodes, links, and facilities) that are undergoing maintenance as down and routes the LSPs traversing through those entities through other paths, thereby ensuring that there is no downtime because of the maintenance.

The maintenance event that you add is listed under the Maintenance tab, where you can view the progress of the maintenance event from its status; see ["About the Maintenance Tab" on page 777](#).

To schedule and to monitor a maintenance event:

1. On the Paragon Automation GUI, navigate to the Topology page (**Network > Maintenance**).
2. Schedule a maintenance event from the Maintenance tab of the Network Information table on the Topology page; see ["Add a Maintenance Event" on page 778](#).

After you schedule the maintenance event, the maintenance event is listed in the Maintenance tab of the Network Information table; see ["Maintenance Event Overview" on page 775](#)

3. Identify the LSPs that are affected by the maintenance event by using the Controller Status of the LSPs in the Tunnels tab of the Network Information table. [Table 30 on page 113](#) lists the Controller Status values for LSPs affected by a maintenance event.
4. After the maintenance event is complete, ensure that the LSPs that were previously affected by the maintenance event are up and working fine:

NOTE: After the maintenance event is completed, the Reoptimize Tunnels upon Completion setting, which you configure while scheduling a maintenance event, determines whether Paragon Pathfinder recalculates the most optimum path or not. For more information, see ["Add a Maintenance Event" on page 778](#).

- a. On the **Tunnels** tab of the Network Information Table, ensure that the Controller Status field for each LSP, previously affected by the maintenance event is empty.
- b. Select the LSPs in the Network Information Table to view the routes taken by the LSPs (on the network topology map).
- c. (Optional) Use the **Diagnostic** tool present on the top-right side of the Network Information Table to:

- Check the details of the LSP, as configured on the router at the head-end of the LSP to confirm that the LSP is configured correctly.
 - Send ping traffic over the LSP to check whether Node A is able to reach Node Z.
 - Send traceroute traffic over the LSP to check the intermediate nodes between Node A and Node Z and if any forwarding issues are present in the LSP.
- d. For each LSP that was previously affected, select the LSP and click **View > Tunnel Traffic** to ensure that traffic is flowing smoothly through the LSP, and click **View > Delay** to ensure that delay is within the acceptable limits.

You have now completed all the tasks that are related to the maintenance event.

Table 30: Controller Status for LSPs under Maintenance

Controller Status Field	Description
Maint_Reroute	The LSP is rerouted due to maintenance.
Maint_NotHandled	The LSP is not a part of the ongoing maintenance event because the LSP is not controlled by Paragon Pathfinder.
Maint_NotReroute_DivPathUp	The LSP is not rerouted due to the maintenance event because a standby path is already up and functioning.
Maint_NotReroute_NodeDown	The LSP is not rerouted because the maintenance event is affecting the endpoints of the LSP.

RELATED DOCUMENTATION

[Maintenance Reports Overview](#) | 918

Reroute LSPs Automatically

Paragon Pathfinder can automatically reroute LSPs based on delay, packet-loss threshold, and link utilization in LSPs . For example, you can configure Paragon Pathfinder to reroute an LSP when the utilization percentage of a link through which the LSP traverses reaches 75 percent of the link capacity. See "[Reroute LSP Overview](#)" on page 718 for more details.

To reroute LSPs, Paragon Pathfinder uses the variables listed in [Table 31 on page 114](#).

Table 31: Parameter for Rerouting LSPs Automatically

Parameter	Description	Range
reroute-minimum-interval	<p>Interval in minutes during which Paragon Pathfinder periodically checks the delay, packet loss, and link utilization in LSPs.</p> <p>Paragon Pathfinder reroutes the LSPs when the measured values for delay, packet loss, or link utilization in an LSP exceeds the corresponding thresholds configured in Paragon Automation. If you do not specify an interval, Paragon Pathfinder does not reroute the LSPs when the measured values exceed the configured thresholds.</p>	1 through 300
Maximum delay	<p>Defines the delay (in milliseconds) at which Paragon Pathfinder must reroute an LSP. For real-time traffic such as VoIP, set the maximum delay to less than 150ms.</p>	
Link utilization threshold (link-utilization-threshold)	<p>Threshold value (in percentage) for link utilization. When traffic on a link exceeds this value, Paragon Pathfinder triggers the rerouting of LSPs. If a threshold is not specified, LSPs are not rerouted. If 0 is specified, links are blocked (no traffic is allowed through the links) and the LSPs are not rerouted.</p>	Range: 0 through 100

Table 31: Parameter for Rerouting LSPs Automatically (*Continued*)

Parameter	Description	Range
Packet loss threshold (packet-loss-threshold)	Threshold value (in percentage) for packet loss on all links. If the packet loss on a link exceeds this value, the link is considered unstable, and the path computation server (PCS) reroutes the LSP and triggers a maintenance event on the link. If you do not specify a value or if you specify 0, the PCS does not reroute the traffic to another LSP in case of packet loss.	Range: 0 through 100

To reroute LSPs automatically:

1. Do one of the following on the Topology page (**Network > Topology**).
 - a. To add a tunnel (LSP), click **Add** icon (+).
The Add Tunnels page appears.
 - b. If you are editing an existing tunnel, click **Edit** icon (pencil).
The Edit Tunnels page appears.
2. On the Constraints tab of the Add Tunnels page or Edit Tunnels page, configure **Maximum Delay**. See ["Add a Single Tunnel" on page 729](#).
3. In the Analytics section of the Path Computation Server settings on the Pathfinder page (**Configuration > Network Settings > Pathfinder Settings**), configure the following parameters:
 - reroute-minimum-interval
 - link-utilization-threshold
 - packet-loss-threshold

See ["Modify Pathfinder Settings From the GUI" on page 206](#).
4. To view LSPs that might have been rerouted:
 - a. Right-click anywhere on the Topology page and select **Timeline**.

The timeline is refreshed every two minutes and displays network events in a chronological order. Rerouted LSPs are displayed when you right-click Link Down events.
 - b. Select the LSP and click **View > Tunnel Traffic** to ensure that traffic is flowing smoothly through the LSP, and click **View > Delay** to ensure that delay is within the acceptable limits.

RELATED DOCUMENTATION

[Modify Pathfinder Settings From the Pathfinder CLI](#) | 198

Add and Check Container LSPs

A container LSP consists of a group of sub-LSPs in which the sub-LSPs are dynamically added or removed based on bandwidth requirements. Paragon Pathfinder allows you to configure container LSPs with a maximum of 32,767 sub-LSPs. For details about container LSPs, see "[Container LSP Overview](#)" on [page 768](#).

Before you add a container LSP, ensure that you have completed the steps to acquire the network topology, as explained in "[Acquire and View the Network Topology](#)" on [page 109](#).

To add and check a container LSP:

1. On the Paragon Automation GUI, navigate to the Topology page (**Network > Maintenance**).
2. Add a container LSP in the Container LSP tab of the Network Information table (on the Topology page); see "[Add a Container LSP](#)" on [page 770](#).

After you add the LSP, the LSP is listed on the Container LSP tab on the Topology page .

3. Verify that the container LSP is up and working fine. In the Container LSP tab:
 - a. Verify that the Operational Status field in the Network Information Table is Active.
 - b. (Optional) Select the LSP in the Network Information Table to view the routes taken by the LSP (on the network topology map).
 - c. (Optional) Use the **Diagnostic** tool present on the top-right side of the Network Information Table to:
 - Check the details of the LSP, as configured on the router at the head-end of the LSP, to confirm that the LSP is configured correctly.
 - Send ping traffic over the LSP to check whether Node A is able to reach Node Z.
 - Send traceroute traffic over the LSP to check the intermediate nodes between Node A and Node Z, and if any forwarding issues are present in the LSP.
 - d. (Optional) Select the LSP and click **View > Tunnel Traffic** to ensure that traffic is flowing smoothly through the LSP, and click **View > Delay** to ensure that delay is within the acceptable limits.
4. Add a container normalization task to increase or decrease the number of sub-LSPs in the container LSP; see "[Add a Container Normalization Task](#)" on [page 975](#).

The normalization task configures the path computation server (PCS) to add sub-LSPs when more bandwidth is required and removes sub-LSPs when less bandwidth is required.

5. Select the LSP and click **View > Tunnel Traffic** to ensure that traffic is flowing smoothly through the LSP, and click **View > Delay** to ensure that delay is within the acceptable limits.

RELATED DOCUMENTATION

| [About the Container LSP Tab | 769](#)

Paragon Planner

IN THIS CHAPTER

- [Obtaining and Importing Network Files | 118](#)
- [Plan Network for Optimum Performance | 119](#)
- [Simulate Failure Scenarios | 120](#)

Obtaining and Importing Network Files

Paragon Planner helps you to simulate various scenarios in a network without affecting the actual network. You can obtain the network files for simulating and analyzing network scenarios in Paragon Planner by running the device collection and network archive tasks in Paragon Pathfinder, and then importing the network files into Paragon Planner. You can then use the imported files to simulate various network scenarios, such as failure of network elements (nodes, links, interfaces, and so on), and see how the failures affect the network.

To obtain network files for simulation and analysis in Paragon Planner:

1. Run the device collection task in Paragon Pathfinder and ensure that you select the **Store Collection for Planner** check box, so that the specification files and raw data from the network devices are available in Paragon Planner. See ["Add a Device Collection Task" on page 978](#).
2. Run the Network archive task in Paragon Pathfinder to create the network model for use in Paragon Planner. See ["Add a Network Archive Task" on page 990](#).
3. Import the network data that you obtained in the preceding step into the Paragon Planner application:
 - a. Access the Paragon Planner Desktop Application. See *Access Paragon Planner Desktop Application*.
 - b. Import the network topology by using the Import Network wizard (**File > Import Network Wizard** on the menu bar). See *Router Data Extraction Overview* The topology of the network appears in the Map window

You can proceed with running a simulation or exit Paragon Planner and run a simulation later. For more information, see:

- ["Plan Network for Optimum Performance" on page 119](#)
- ["Simulate Failure Scenarios" on page 120](#)

RELATED DOCUMENTATION

| [Key Paragon Planner Features](#)

Plan Network for Optimum Performance

You can simulate a network in Paragon Planner and use the simulation to check the network's performance when you add, modify, or remove (delete) network elements from the network. The network elements that you can add, modify, or remove include nodes, links, tunnels, demands, and so on.

To plan your network for optimum performance:

1. Open the network from the Network Browser window (**File > Open Network Browser**).

The topology map of the network appears on the Map window.

2. Analyze the network topology and the traffic patterns in the topology map; see *Topology Window Overview*.

3. Add, delete, or modify the network elements in your network to view the impact on the network performance:

- a. Click **Network** on the main menu.

The Network Info section is displayed below the topology map. The Network Info section lists all the network elements (nodes, links, interfaces, and so on) in different tabs.

- b. Do one or more of the following:

- To add a network element, click the tab for that network element and click the **Add** button.

The Add *Element-Name* window (for example, Add Interface) appears.

- To modify a network element, click the tab for that network element and click the **Modify** button.

A page to modify the selected network element appears.

- To delete a network element, click the tab for that network element and click the **Delete** button.

You are asked to confirm the delete action.

For more information on the add, modify, and delete actions, see the *Network Menu* chapter of the *Paragon Planner Desktop Application User Guide*.

4. View the changes on the topology map in link utilization, congestion, and traffic flow (demands) in the network because of the changes you made.
5. Repeat steps 3 and 4 to modify the topology as needed until you get the desired performance.
6. Save the network files for future analysis; see *Saving a File*.

RELATED DOCUMENTATION

| *Sizing Tunnels and Demands*

Simulate Failure Scenarios

Paragon Planner allows you to simulate failure scenarios on a network to determine the resiliency of the network; see *Simulation Menu Overview*.

You can perform two types of simulations:

- By using pre-defined scenarios. Use this type of simulation to analyze network performance by failing:

- One network element type (node, link, site, linecard, and so on) at a time (single failure).

For example, if you select to fail nodes, failure scenarios are executed by failing all nodes in the network, one node at a time.

- Two or three types of network elements at a time (multiple failures).

When you select multiple failures, failure scenarios are executed by failing two or three types of selected network elements randomly. The scenarios are executed for the number of times that you specify.

- By manually selecting the network elements (interactive scenarios) to fail. In this type of simulation, you select specific network elements of a particular type (for example, nodes, links, interfaces, and so on) or a combination of different types of network elements to fail.

See *Simulation Menu: Predefined and Interactive Scenarios* for details.

To simulate failure scenarios:

1. Open the network in which you want to simulate the failure scenarios, from the Network Browser window (**File > Open Network Browser**).

The topology map of the network appears on the Map window.

2. On the main menu, select **Tools > Options > Failure Simulation** to review the Failure Simulation Options and if needed, modify the options. See *Failure Simulation Options* section of the *Tools: Options Menu*.

3. To simulate failure scenarios on the network, do one of the following:

- For interactive failure simulations:
 - a. Select **Simulation > Interactive Scenarios...** on the menu bar.

The Interactive Scenarios... window appears.

- b. Select the network elements to fail.
- c. Go to step 4.

- For predefined failure simulations:
 - a. Select **Simulation > Simulation Scenarios** on the menu bar.

The Simulation Scenarios window appears.

- b. Configure the simulation options explained in [Table 32 on page 121](#).

4. Click **Run**.

Paragon Planner triggers the failure simulation and after the simulation is complete, a message appears indicating that a failure simulation report is generated.

5. View the failure simulation reports on the Reports window (**Report > Report Manager > Simulation Report**). See *Report Manager: Simulation Reports* for details.

6. Save the network in Paragon Planner for future analysis; see *Saving a File*.

Table 32: Parameters to Configure for Predefined Simulation Scenarios

Tab	Action/Description
Single Failure	Select a network element type (Node, Card, Link, Site, Slot, SRLG [Shared Risk Link Group]/Facility, or Parallel links) to fail one network element at a time.

Table 32: Parameters to Configure for Predefined Simulation Scenarios (*Continued*)

Tab	Action/Description
Multiple Failures	<p>Do the following:</p> <ul style="list-style-type: none"> • From the Multiple Network Element Failure field: <ul style="list-style-type: none"> • Select Double to run simulation by failing two types of network elements at the same time. • Select Triple to run simulation by failing three types of network elements at the same time. • Click Cases to Simulate and enter the number of failures that you want to simulate. For example, if you select this option and enter 10 in the text field, Paragon Planner simulates the network ten times by failing a random combination of the selected network element types. • Click Exhaustive Failure to run simulation by failing all combinations of the chosen network element types. For example, if you chose to run an exhaustive failure simulation by failing two network elements—links and nodes— at a time. Paragon Planner fails a combination of two nodes, a node and link, and two links to simulate network failures. • Under Failure Types, select the types of network elements that you want to fail—Node, Link, Site, Slot, Card, SRLG/Facility, Parallel Links.
Report Options	<p>Select the types of reports that you want to generate and enter an alphanumeric value in the Report Runcode text box.</p> <p>The Report Runcode is an extension that Paragon Planner adds to the filenames of the reports that it generates. Runcodes help you to distinguish between reports that are generated by using the same data set in different sessions.</p> <p>See the <i>Report Options</i> section of <i>Simulation Menu: Predefined and Interactive Scenarios</i>.</p>

Table 32: Parameters to Configure for Predefined Simulation Scenarios (*Continued*)

Tab	Action/Description
Advanced Options	<p>Do the following:</p> <ul style="list-style-type: none"> • Select the fast reroute (FRR) mode to be used during the simulation— Normal, FRR only, or Normal + FRR. See <i>Failure Simulation Options</i> section of <i>Tools: Options Menu</i> for details. • Under Replay Up-Down Sequence, click Browse... and select a custom simulation script to be run, instead of running the default simulation scripts provided by Paragon Planner. • From the Simulation field, select one of the following: <ul style="list-style-type: none"> • Every Event—A network element is brought up or down for every event (<i>network-element</i> up or <i>network-element</i> down) as indicated in the custom script. • When WAIT/RESET is Encountered—A network element is brought up or down based on a WAIT or RESET command present in the custom script: <ul style="list-style-type: none"> • WAIT—The network elements are retained in their current state (up or down) and traffic that was rerouted due to network elements going down are also retained in the new paths. • RESET—The network elements are reset to a state that they were present before the start of the simulation and traffic that was rerouted due to network elements going down are also rerouted through their original paths. <p>See the <i>Advanced Options</i> section of <i>Simulation Menu: Predefined and Interactive Scenarios</i></p>

RELATED DOCUMENTATION

| *Perform Failure Simulation and Assess the Impact*

Paragon Insights

IN THIS CHAPTER

- [Monitor and Troubleshoot Device and Network Health | 124](#)
- [Analyze Root Cause of Network, Device, and Service Issues | 125](#)
- [Bring Your Own Ingest Default Plug-in Workflow | 127](#)
- [Bring Your Own Ingest Custom Plug-in Workflow | 127](#)
- [Generate and View Health Reports | 128](#)

Monitor and Troubleshoot Device and Network Health

Paragon Insights gathers and analyzes configuration and telemetry data from your network devices. You can view this data and monitor the status of your network in real-time by using the Paragon Automation UI.

To view and monitor the status of your network devices, and troubleshoot issues:

1. View statistics and status information of devices (vendor-wise), device group health, list of networks and their health, and status of network groups. This information is displayed in a carousel as widgets on the Dashboard page. For more information, see ["About the Dashboard Page" on page 38](#).
2. Create and view charts, graphs, and other visuals to help you organize and understand device and network data by using:
 - Charts, which is an in-built data visualization tool that you can use to create time series graphs, histograms, and heatmaps. For more information, see ["Charts Overview" on page 401](#) and ["Manage the Charts Page" on page 404](#).
 - Grafana, which is an open-source data visualization tool that you can use to create and view visuals, view pre-populated graphs, and run queries by using syntax parameters. For more information, see ["Grafana Overview" on page 423](#) and ["Monitor Network Device Health Using Grafana" on page 424](#).

3. Generate health reports to view information on device and network health, device-related hardware and software specifications, and alarm statistics. For more information, see ["Generate and View Health Reports" on page 128](#).
4. To troubleshoot device and network issues, you can use built-in diagnostics tools to:
 - Diagnose problems after initial setup and validate the basic functionality of Paragon Insights. For more information, see ["Use the Self Test Tool" on page 564](#).
 - Determine whether the devices that you added to your network are up and reachable. For more information, see ["Use the Reachability Test" on page 566](#).
 - Verify that the data ingest methods, such as OpenConfig, iAgent, and SNMP, are applied on network devices to ensure data collection and monitor the flow of traffic. For more information, see ["Use the Ingest Test Tool" on page 567](#).
 - Identify why a device or a rule displays a "no-data" status. For more information, see ["Use the No-Data Tool" on page 568](#).

Analyze Root Cause of Network, Device, and Service Issues

You can use Paragon Automation to identify the root cause of issues related to network, device, or service health. Issues are logged as error events, which are then displayed as alarms in the Paragon Automation GUI. Paragon Automation performs root cause analysis (RCA) and combines these alarms and identifies the original cause of the issue. The original cause is displayed in a smart alarm along with the resulting dependent alarms.

Smart alarms allow you to focus on the core reason for the issue and you don't have to go through multiple alarms to find the root cause. After you identify the root cause, you can remediate the issue as needed. For more information about root cause analysis, see ["Understand Root Cause Analysis" on page 369](#).

NOTE: Before you perform RCA, ensure that you have onboarded devices, created device groups, and configured topology acquisition, so that Paragon Automation can acquire the network topology. For more information, see ["Acquire and View the Network Topology" on page 109](#).

To analyze the root cause of network, device, or service issues, you must:

1. Create one or more rules.

Rules define how to collect and analyze telemetry data about the network and generate notifications when anomalies occur. You can create rules to monitor network, device, and service health, network

resource key performance indicators (KPIs), various system parameters, and so on. For more information, see ["Configure a Custom Rule in Paragon Automation GUI" on page 343](#).

2. Create and run a playbook instance.

Playbooks are collections of rules for addressing specific use cases and run on device or network groups. You must include the rules added in Step 1 to playbooks in order to deploy them to network devices. Alarms are generated based on the configured playbooks and rules. For more information, see ["Create a Playbook Using the Paragon Insights GUI" on page 309](#).

3. Add one or more resources.

Resources are network elements such as devices, interfaces, protocols, label switched paths (LSPs), IPSec tunnels, and so on. To identify the root cause of a network device, or service issue, you must map resources across rules. For more information, see ["Add Resources for Root Cause Analysis" on page 375](#).

4. Configure dependencies between resources.

Multiple alarms on different resources can be traced back to the original issue, only when dependencies are configured between the resources. When you configure dependencies between resources, Paragon Automation can combine a set of error events on those resources and identify the root cause of the errors. For more information, see ["Configure Dependency Between Resources" on page 378](#).

5. View smart alarms and the root cause.

If a network, device, or service issue occurs, alarms are generated based on the playbooks running in the network and on the network devices. When KPIs cross predefined thresholds, causing multiple KPI anomalies, multiple alarms are generated.

Paragon Automation categorizes all linked alarms into a single group called a smart alarm and identifies the root cause of the sequence of errors. The alarms are displayed in a collapsible tree structure with the root cause of the alarms displayed at the top of the tree. The alarms are listed underneath the root cause in the order that the alarms occurred, thereby displaying the different issues triggered by the root cause. You can view smart alarms on the Alerts page (**Monitoring > Alarms and Alerts > Alerts**). For more information, see ["About the Alerts Page" on page 850](#).

After you analyze the smart alarms, you can remediate the issue as needed.

NOTE: The RCA feature is disabled, by default, in Paragon Automation Release 23.1. For instructions on how to enable the feature, see the [Release Notes](#).

RELATED DOCUMENTATION

| No Link Title

Bring Your Own Ingest Default Plug-in Workflow

You can use default plugins to measure metrics that are unique to your network. To load a default plug-in, you require at least Paragon Automation Release 21.3. After you load the default plug-in, it is listed in the Default Plugins tab of the BYOI Plugins page. All default plug-ins must be mapped to existing or new device groups. Ensure that you deploy playbook instances on the device groups mapped to the plug-in ingest. Custom plug-ins do not support the default rules, playbooks, and device groups in Paragon Automation. Therefore, you should write your own rules and playbooks and add a device group for the custom ingest plugin.

The following tasks comprise the end-to-end workflow to use the BYOI default plug-in:

1. Load the default plug-in to the Paragon Insights server. See ["Load BYOI Default Plug-ins" on page 541](#).
2. Create an instance of the plug-in you earlier loaded. See ["Configure Bring Your Own Ingest Default Plug-in Instances" on page 542](#).
3. Map the plug-in instance to sensors and device groups that can then use the plug-in. See ["Configure Ingest Mapping for Default BYOI Plug-in Instances" on page 558](#).
4. Add the default plug-in as a sensor in a rule. See ["Configure a Custom Rule in Paragon Automation GUI" on page 343](#).
5. Add the rule to a playbook. See ["Create a Playbook Using the Paragon Insights GUI" on page 309](#).
6. Deploy a playbook instance. See ["Manage Playbook Instances" on page 312](#).

RELATED DOCUMENTATION

[Understand Bring Your Own Ingest | 540](#)

Bring Your Own Ingest Custom Plug-in Workflow

You can use custom plug-in when you want to stream pre-existing telemetry data to Paragon Automation for analysis. To load a custom plug-in, you require at least Paragon Automation Release 22.1. After you successfully load the custom plug-in, it is listed in the Custom Plugins tab of the Bring Your Own Ingest Plugin page.

You must perform the following tasks to use custom BYOI plug-in in Paragon Automation:

1. Custom BYOI plug-in requires an ingest image file and a Kubernetes YAML file. Build the ingest image and load the image and YAML file. See ["Build and Load BYOI Custom Plug-in Images" on page 544](#).
2. Create an instance of the plug-in you earlier loaded. See ["Configure Bring Your Own Ingest Custom Plug-in Instances" on page 555](#).
3. Create a device group for the custom plug-in. See ["Add a Device Group" on page 177](#).
4. Create rules and playbooks that use the custom plug-in. ["Use Sample Rule and Playbook Configurations for BYOI Custom Plug-in Instances" on page 557](#).

RELATED DOCUMENTATION

[Understand Bring Your Own Ingest](#) | 540

Generate and View Health Reports

Paragon Insights gathers and analyzes configuration and telemetry data from your network devices. You can view this data in real-time, and also generate health reports by using the Paragon Automation UI. Health reports provide information on device and network health, device-related hardware and software specifications, and alarm statistics.

To generate and view health reports:

1. Configure a report schedule.

You can set the date, time, duration, type, and number of times a report is generated from the Scheduler Settings page. For more information, see "[Configure Scheduler Settings](#)" on page 630.

2. Set report destination.

You can configure settings to either send a copy of a health report over email or save a copy on the server from the Destination Settings page. For more information, see "[Configure Destination Settings](#)" on page 633.

3. Configure report settings.

You can configure report settings to generate a report on-demand, determine the report format, and so on. For more information, see "[Configure Report Settings](#)" on page 628.

4. View and download the latest health reports from the Health Reports page. You can download health reports only if the destination type for a report is set to disk. For more information, see "[View and Download Health Reports](#)" on page 912.

5. (Optional) Compare the differences between two health reports. For more information, see "[Compare Differences in Health Reports](#)" on page 913.

4

PART

Manage Devices and Network

[Devices](#) | 130

[Device Groups](#) | 175

[Device Images](#) | 188

[Network](#) | 197

[Network Groups](#) | 264

[Topology Filter](#) | 270

Devices

IN THIS CHAPTER

- [Devices Overview | 130](#)
- [Zero-Touch Provisioning Overview | 132](#)
- [Configure a DHCP Relay for ZTP | 134](#)
- [Add Devices Overview | 135](#)
- [About the Devices Page | 138](#)
- [Supported Devices | 144](#)
- [Add Devices | 149](#)
- [Upgrade the Device Image | 156](#)
- [View Device Statistics and Inventory information | 157](#)
- [View and Manage Device Configuration | 165](#)
- [Edit Devices | 168](#)
- [Delete Devices | 173](#)

Devices Overview

You can manage the devices in your network from the Devices (**Configuration > Devices**) page. Paragon Automation supports Zero-Touch Provisioning (ZTP), which you can use to provision the device while onboarding the device. For information about ZTP, see ["Zero-Touch Provisioning Overview" on page 132](#).

You can use Paragon Automation to manage both Juniper Networks devices and third-party devices (for example, Cisco IOS XR devices). A managed device refers to a device that the base component of Paragon Automation can discover in a network, and therefore can configure or monitor. An unmanaged device refers to a device that cannot be discovered, configured, or monitored by the base component of Paragon Automation.

[Table 33 on page 131](#) lists the tasks that you can perform on the Juniper Networks, Nokia, and Cisco IOS-XR devices in this release.

Table 33: Tasks you can perform for Juniper Networks, Nokia, and Cisco IOS XR Devices

Task	Juniper Networks Devices	Nokia Devices	Cisco IOS XR Devices
View Inventory	Yes	Yes NOTE: Chassis view is not supported. You cannot view the features supported or software installed.	Yes NOTE: Chassis view is not supported. You can view the interfaces, features supported, and software installed.
Automatic Resynchronization with Network	Yes	No	Yes
View Active Configuration	Yes NOTE: You can view the active device configuration in the SET, Junos Native, and XML formats.	Yes NOTE: You can view the active device configuration in the Set and XML formats.	Yes NOTE: You can view the active device configuration in the Set and XML formats.
Configuration Versions- View configuration version history, rollback to a previous configuration version, compare different configuration versions, or edit a device configuration	Yes	Yes	Yes
Backup and restore device configuration	Yes	Yes	Only backup of device configuration is supported.
Reboot Device	Yes	Yes	Yes

Table 33: Tasks you can perform for Juniper Networks, Nokia, and Cisco IOS XR Devices (Continued)

Task	Juniper Networks Devices	Nokia Devices	Cisco IOS XR Devices
Upgrade image	Yes	No	No NOTE: Device image must be upgraded manually.
ZTP	Yes	No	No

For a list of all the devices that are supported by Paragon Automation, see ["Supported Devices" on page 144](#).

RELATED DOCUMENTATION

[Configuration Templates Overview | 278](#)

[Device Templates Overview | 298](#)

Zero-Touch Provisioning Overview

IN THIS SECTION

- [Benefits | 133](#)

Zero-Touch Provisioning (ZTP) allows you to provision new Juniper Networks devices in your network automatically, with minimal manual intervention. The ZTP feature uses Dynamic Host Configuration Protocol (DHCP) to provision devices. You can use ZTP to onboard more than one device at a time.

NOTE:

- In Paragon Automation, you can onboard only MX Series and QFX Series devices by using ZTP.
- For ZTP on Junos OS devices, the supported versions are Junos OS Release 21.1 or later; see ["Deploy an Image" on page 194](#).
- On devices running Junos OS Evolved, the ZTP process applies the bootstrap configuration and leases IP addresses to the device. ZTP does not upgrade the image on the device.
- To use ZTP for onboarding devices that are present on a subnet that is different from the subnet in which Paragon Automation is installed, see ["Configure a DHCP Relay for ZTP" on page 134](#).

To use ZTP for onboarding a device that is present on the same subnet as Paragon Automation involves the following steps:

1. Provide the inputs for device onboarding, such as the available IP addresses (range) to be leased, corresponding gateway information, and device details like root password and serial number, see ["About the Devices Page" on page 138](#).

Paragon Automation triggers a job and the job lists the progress of the ZTP process.

2. A DHCP server, previously configured in Paragon Automation, leases an available IP address to the device requesting for an IP address.
3. The device uses the IP address received from the DHCP server to connect with Paragon Automation.
4. The bootstrap configuration and software images required for ZTP (served through the internal HTTP server of Paragon Automation) are installed on the device.
5. The device establishes a NETCONF session with Paragon Automaton to synchronize the inventory, configurations, and the configuration version.
6. (Optional) You can monitor the progress of the ZTP job on the Jobs (**Monitoring > Jobs**) page.
7. After the ZTP job is completed, you can navigate to the Devices (**Configuration > Devices**) page to view the list of onboarded devices.

Benefits

ZTP in Paragon Automation offers the following benefits:

- Simplified, faster, and automated deployment of device configurations.
- Auto-generated device configurations that are more accurate.

- Faster provisioning of devices with minimal manual interventions.

RELATED DOCUMENTATION

[Upgrade the Device Image | 156](#)


[Configuration Templates Overview | 278](#)

[Device Templates Overview | 298](#)

Configure a DHCP Relay for ZTP

To use ZTP for onboarding devices that are present on a subnet that is different from the subnet in which Paragon Automation is installed, you must configure a DHCP relay, as shown in [Figure 8 on page 134](#).

Figure 8: Sample DHCP Relay Configuration for ZTP



The DHCP relay can be an MX Series device or a Linux-based (CentOS) Virtual Machine (VM). For information about using an MX Series device as a DHCP relay, see [DHCP Relay Agent](#).

To configure a DHCP relay in a CentOS-based VM:

1. Log in to the VM as a root user.
2. Install the DHCP relay package on the VM.

```
root@host# yum install dhcp
```

NOTE: The DHCP relay package is usually available by default in a CentOS-based VM. However, if the package is not available, the `yum install dhcp` command fetches the package and installs it.

- Execute the following command to run the DHCP relay service.

```
root@host# dhcrelay -4 -d -i <interface-name> <dhcp-service-external-ip>
```

where:

interface-name is the interface on the VM that is connected to the device to be onboarded to Paragon Automation.

dhcp-service-external-ip is the Paragon Insights services virtual IP (VIP) address that was provided during the installation.

You can also obtain the Paragon Insights services VIP address by running the following command on the Paragon Automation primary node:

```
root@device# kubectl get svc -n ems | grep -i ztpservicedhcp
NAME                                TYPE             CLUSTER-IP      EXTERNAL-IP      PORT(S)          AGE
ztpservicedhcp                      LoadBalancer     10.102.3.100    192.168.15.155  67:30241/UDP     37h
```

The IP address listed under EXTERNAL-IP is the Paragon Insights services VIP address.

You can now use the VM as a DHCP relay to connect Paragon Automation and the devices to be onboarded to Paragon Automation.

RELATED DOCUMENTATION

[Zero-Touch Provisioning Overview](#) | 132

Add Devices Overview

You use the *Add Device* option to discover and synchronize all devices connected to your network. Paragon Automation classifies devices as managed and unmanaged.

- Managed devices refer to devices that Paragon Automation discovers by using hostname or IP address. Paragon Automation synchronizes configuration of managed devices by using NETCONF.

For a complete list of managed devices, see ["Supported Devices" on page 144](#).

- Unmanaged devices refer to devices that third-party applications (for example, Anuta ATOM) discover and that Paragon Automation supports. Paragon Automation cannot synchronize configuration of unmanaged devices.

Paragon Automation lists the hostname and IP address of the unmanaged devices on the Devices page.

To discover and onboard a device by using Paragon Automation, you must provide the IP addresses, hostnames, and IP subnets to be discovered and the credentials to connect to the devices.

NOTE:

- Netconf over ssh should be enabled on Cisco IOS XR devices and you should have the privilege to manage Netconf on the devices.
- For devices with dual Routing Engines, Chassis Cluster, or Virtual Chassis, enter the Virtual IP address (VIP) or the IP address of the primary device
- The values that you enter to specify the targets, and credentials are persistent from one discovery operation to the next, so you do not have to reenter information that is the same from one operation to the next

You can add devices in a network to Paragon Automation in one of the following ways:

- Discover the devices already present in a network by providing the device credentials.

You can use this option to discover devices that are already present in your network. You can discover the devices by providing a list of IPv4 addresses, subnet addresses, IP ranges or hostnames.

You can also use a CSV file to provide the targets to discover the existing devices in a network.

See ["Discover Devices" on page 149](#) for more information about discovering devices.

- Add devices manually by specifying their details such as serial number and device model and the root password.

You can use this option to add a specific set of devices; for example, add a set of new EX Series devices. When you add a new device, the device is automatically provisioned by zero-touch provisioning (ZTP). For details on ZTP, see ["Zero-Touch Provisioning Overview" on page 132](#)

You can also use a CSV file to specify the details of the devices to be added.

See ["Add New Devices" on page 153](#) for more information about adding new devices.

When you add a device, Paragon Automation executes the following tasks:

1. Paragon establishes connection with the devices based on the specified device targets.
2. Deploys the default configuration templates included in the device template associated with the device-family to which the device belongs.
3. Deploys telemetry configuration for monitoring.
4. Captures and stores the device inventory data.
5. Captures and stores the device configuration.
6. Applies any additional configuration template, for example for configuring BGP, if the configuration template is assigned for the device.
7. Synchronizes the device configuration in the database.
8. Pins the configuration.

The time required to complete the discovery process or adding the devices depends on multiple factors such as the number of devices you are discovering, the size of configuration and inventory data on the devices, the network bandwidth available between Paragon Automation and the devices, and so forth.

After your devices are successfully discovered or added, you can view them on the Devices page (Configuration > Devices) of Paragon Automation.

- For managed devices, after the discovery, the Management Status should display “Up” and the Sync Status “In Sync”, which indicates that Paragon Automation has established a session with the devices and that the configuration and inventory data in Paragon Automation is in sync with the data on the devices.
- For unmanaged devices, the Management Status should display as “Unmanaged” and the Sync Status “Unknown”, which indicates that Paragon Automation has added the device details to its database.

RELATED DOCUMENTATION

[Upgrade the Device Image | 156](#)

[View Device Statistics and Inventory information | 157](#)

[View and Manage Device Configuration | 165](#)

About the Devices Page

IN THIS SECTION

- [Tasks You Can Perform | 138](#)
- [Field Descriptions | 141](#)

The Devices page in Paragon Automation helps you manage the devices in your network. To access the Devices page, click **Configuration > Devices**.

Tasks You Can Perform

Users with the SP Administrator role can perform the following tasks from this page, while users with the SP Operator role have read only capabilities.

- View the list of devices managed by Paragon Automation, and their details.

To view details of a specific device, select the device and click **More > Detail**. Alternatively, hover over the device name and click the **Details** icon that appears.

The Device Details pane appears on the right side of the page, displaying basic information, such as host name, IP address, device family, OS version, platform, serial number, sync status, and management status, and synchronization status of the device. Click the **close** (x) icon to close the pane.

- Add new devices or discover devices existing in the network topology; see ["Add Devices" on page 149](#).
- View details of a device's inventory.

To view details of a device's inventory, click the device link or select the device and click **More > View Inventory**. The *Device-Name* page appears, displaying the details on the Inventory tab of the Devices page.

- Synchronize the device inventory details with the network

NOTE:

- Juniper Networks devices are automatically synchronized with the network whenever a change in inventory, such as change in configuration, interface link or admin status, or FRU insertion or deletion, is detected in the system logs.
- Device discovery and resync of inventory takes more time if the configuration on the device is large (generally over 4 MB).

To synchronize the device inventory details of a device, select the device and click **More > Resynchronize With Network**. A job is triggered to synchronize the device inventory stored in database with the network.

Use this option to synchronize the inventory of a device when its **Sync Status** is *Unknown* or *Out-of-Sync*.

- View the current configuration used by the device (active configuration)

To view the active configuration of a device, select the device and click **More > View Active Configuration**. The Active Configuration for *Device-Name* page appears, displaying the configuration in the Set, Junos Native, and XML formats for Juniper Networks devices and in the Set and XML formats for CISCO-IOS XR devices.

- View the different configuration versions used on a device

To view the different configuration versions used on a device, select a device and click **More > Configuration Versions**. The Configuration Versions for *Device-Name* page appears, displaying the current configuration and the version history of the different configurations used on the device.

On the Configuration Versions for *Device-Name* page, you can perform the following tasks:

- Compare two versions of a configuration—Select the configurations to be compared and click **Compare**. The Compare Configuration page appears, displaying the configurations side-by-side for you to view and compare.
- Rollback to a previous configuration—Select the configuration that you want to rollback to and click **Rollback**. The device configuration is rolled back to the selected configuration version.
- Pin or unpin a device configuration—Select the configuration and click the **Pin** icon. The device configuration is pinned and a small pin icon appears next to the configuration.

Paragon Automation archives only ten configurations. When the next configuration is to be archived, the oldest configuration is deleted. When you pin a configuration, the pinned configuration is not deleted, but the next oldest configuration is deleted.

To unpin a configuration, select the configuration and click the **Unpin** icon. The configuration is unpinned and the small pin icon, present next to the configuration, is removed.

- Add a description for a configuration version—Select the configuration version and click the **Edit** icon. The Add Description page appears, where you can enter the description in the **Description** text box.

- Reboot a device

To reboot one or more devices, select the devices and click **More > Reboot**. A confirmation message appears. Click **Yes**. A message indicating that the reboot has started is displayed along with the job ID. Click the job ID to view the progress of the reboot job in the Jobs page (**Monitoring > Jobs**).

NOTE: Reboot may affect traffic flowing through the device. For the impact of reboot on a device, refer to the respective device documentation.

- Upgrade images on one or more devices; see ["Upgrade the Device Image" on page 156](#).
- Export devices to a CSV file

To export details of all devices to a CSV file, click **More > Export Devices**. A comma-separated values (CSV) file is generated that you can download to your local system.

The CSV file includes information about the device name, platform, OS version, synchronization state, IP address, family to which the device belongs, management status, serial number, and the unique identifier of the device.

- Export inventory details of a device to a zip file

To export the inventory details of a device to a zip file, select the device and click **More > Export Inventory**. A zip file is generated that you can download to your local system.

The zip file includes five different comma-separated value files (chassis.csv, feature.csv, interface.csv, license.csv, and software.csv) that contain complete inventory details.

NOTE: If any of the inventory resources data is empty, the corresponding CSV file is not included in the zip file.

- Edit parameters configured for devices; see ["Edit Devices" on page 168](#).
- Delete devices from Paragon Automation; see ["Delete Devices" on page 173](#)
- Filter Entries—Filter the table entries by adding new filtering criteria.

Hover over the **Filter** (funnel) icon and select **Add Filter**. On the **Add Criteria** page, select the filtering criteria from the **Field** and **Condition** list, and enter the text to be compared in the **Value** field. Click **Add**.

The filtered table entries are listed and the filter criteria name is displayed above the table column names.

You can click the cross (X) icon (next to the filter name) to remove the filtering criteria.

NOTE:

- You can add multiple filtering criteria. Once you add the multiple filtering criteria, you can select the **And** condition to display the entries matching all the filtering criteria or select the **Or** condition to display the entries matching any one of the filtering criteria.
- Quick filter: Save the filtering criteria as quick filters. Once you have added all the filtering criteria, you can save a particular criteria or multiple criteria for future use by clicking **Save**.

On the **Save Filter** window, 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 a specific column in the table.

Hover over the **More Options** (vertical ellipsis) > **Show/Hide Columns** and select the *Column-Name* check box of the columns you want to display in the table.

- Reset Preference—Reset the displayed columns to the default set of columns for each tab in the table.

Hover over the **More Options** (vertical ellipsis) and select **Reset Preference**.

- Sort Entries—Click the column name to highlight the up and down arrows next to the column name. Sort the table entries in ascending or descending order of that column by clicking the up or down arrow respectively.

Field Descriptions

[Table 34 on page 141](#) describes the fields on the Devices page.

Table 34: Fields on the Devices Page

Field	Description
Host Name	The hostname of the device.

Table 34: Fields on the Devices Page (Continued)

Field	Description
IP Address	The IPv4 address of the device.
Platform	The device model. For example, MX960, SRX3500.
Vendor	The name of the device vendor.
Family	The device family to which the device belongs. For example, Juniper-MX, Juniper-SRX
Management IP	The management IP address (IPv4) of the device, which is used for management access.
OS Version	The version of OS that is currently installed on the device.
Serial Number	The serial number of the device.

Table 34: Fields on the Devices Page *(Continued)*

Field	Description
Management Status	<p>The management Status of the device:</p> <ul style="list-style-type: none"> • Device Detected—The device is detected in the network • Discovery Not Initiated—Indicates that discovery is not initiated on the device. • Discovery In Progress—Indicates that the device details (the device parameters, configuration, inventory, and so on) are being captured. • ZTP Not Initiated—Indicates that the ZTP of the device is not yet initiated. • Discovery Failed—Device discovery failed. • Unmanaged—Device is an unmanaged device; that is, the device is not managed by the base component of Paragon Automation. • Maintenance—Device is under maintenance. • Up—Indicates that the connection between Paragon Automation and the device is up. • Down—Indicates that the connection between Paragon Automation and the device is down.
Sync Status	<p>Synchronization status of the device information stored in the Paragon Automation database with the network.</p> <ul style="list-style-type: none"> • In-Sync—The inventory information in database is synchronized with the device inventory in the network. • Sync in Progress—The inventory information in the database is being updated to reflect the changes in the network. • Out-of-Sync—The inventory information in the database is not synchronized with the network. • Unknown—The device is unmanaged or Paragon Automation is unable to connect with the device.

Table 34: Fields on the Devices Page (Continued)

Field	Description
PCEP IP	PCEP IP address used by the device to connect to Paragon Automation for managing LSPs. The PCEP IP address is usually the management IP address of the device.

RELATED DOCUMENTATION

[Zero-Touch Provisioning Overview | 132](#)

[Devices Overview | 130](#)

Supported Devices

[Table 35 on page 145](#) lists all the devices supported by Paragon Automation for the following basic device functions:

- Device discovery
- Device lifecycle management
- Device inventory management
- Device upgrade
- Configuration management through configuration templates

Paragon Automation supports Juniper Networks, Cisco IOS XR, and Nokia devices. For a Juniper device, follow the instructions in the hardware documentation to unbox the device, mount it on a rack, and power on the device. For details about installing a device, see the device's Hardware Guide in the [TechLibrary](#) or the device's [Quick Start Guide](#). Search for the device in the search box provided or navigate to **Routing > View More**, **Switching > View More**, or **Security > View More**.

Table 35: Devices Supported

Device Family	Device Series	Device Models
ACX Series	ACX710	ACX710
	ACX1000	ACX1000
	ACX1100	ACX1100
	ACX2100	ACX2100
	ACX2200	ACX2200
EX Series	ACX5448 ACX7100	ACX5448-D, ACX5448-M, ACX5448-ac-afi, ACX5448-dc-afi, ACX5448-ac-afo, ACX5448-dc-afo ACX7100-48L
	EX2300	EX2300-c-12t, EX2300-c-12p, EX2300-24t, EX2300-24p, EX2300-48t, EX2300-48p, EX2300-48mp
	EX3400	EX3400-24t, EX3400-24t-dc, EX3400-24p, EX3400-48t, EX3400-48t-afi, EX3400-48p
	EX4300	EX4300-24t, EX4300-24t-s, EX4300-24p, EX4300-24p-s, EX4300-32f, EX4300-32f-s, EX4300-32f-dc, EX4300-48t, EX4300-48t-s, EX4300-48t-afi, EX4300-48t-dc, EX4300-48t-dc-afi, EX4300-48p, EX4300-48p-s, EX4300-48mp, EX4300-48mp-s
EX4600	EX4600-40f-afi, EX4600-40f-afo, EX4600-40f-dc-afi, EX4600-40f-dc-afo, EX4600-40f-s, EX4600-40f	

Table 35: Devices Supported (Continued)

Device Family	Device Series	Device Models
	EX4650	EX4650
	EX9204	EX9204
	EX9208	EX9208
	EX9214	EX9214
MX Series	MX204	MX204
	MX240	MX240
	MX480	MX480
	MX960	MX960
	MX2000	MX2000
	MX2010	MX2010
	MX2020	MX2020
	MX10008	MX10008
PTX Series	PTX10001	PTX10001, PTX10001-m20c, PTX10001-36MR
	PTX10003	PTX10003-80c, PTX10003-160c
	PTX10004	PTX10004
	PTX10008	PTX10008
	PTX10016	PTX10016

Table 35: Devices Supported (Continued)

Device Family	Device Series	Device Models
QFX Series	QFX5000	QFX5000
	QFX5110	QFX5110-32q, QFX5110-48s-4c
	QFX5120	QFX5120-32c, QFX5120-48t, QFX5120-48y, QFX5120-48y-8c, QFX5120-48t-6c, QFX5120-32c-afi, QFX5120-32c-afi-t, QFX5120-32c-afo, QFX5120-32c-afo-t, QFX5120-32c-dc-afi, QFX5120-32c-dc-afo, QFX5120-48y-afi, QFX5120-48y-afi2, QFX5120-48y-afi-t, QFX5120-48y-afo, QFX5120-48y-afo2, QFX5120-48y-afo-t, QFX5120-48y-d-afi2, QFX5120-48y-d-afo2, QFX5120-48y-dc-afi, QFX5120-48y-dc-afo, QFX5120-48t-afi, QFX5120-48t-afo, QFX5120-48t-dc-afi, QFX5120-48t-dc-afo, QFX5120-48t-chas
	QFX5130	QFX5130-32cd
	QFX5200	QFX5200-48y, QFX5200-32c-32q
QFX5210	QFX5210-64c	
	QFX5220	QFX5220-32cd, QFX5220-128c, QFX5220-128c-afo, QFX5220-128c-d-afo, QFX5220-32cd-afi, QFX5220-32cd-afo, QFX5220-32cd-d-afi, QFX5220-32cd-d-afo
	QFX10002	QFX10002-72q, QFX10002-36q, QFX10002-60c, QFX10002-76q, QFX10008, QFX10016
	QFX10003	QFX10003-80c, QFX10003-160c

Table 35: Devices Supported (*Continued*)

Device Family	Device Series	Device Models
SRX Series	SRX300	SRX300-poe-ac
	SRX320	SRX320, SRX320-lem, SRX320-poe, SRX320-poe-lem, SRX320-lte-aa, SRX320-lte-ae, SRX320-poe-lte-aa, SRX320-poe-lte-ae
	SRX340	SRX340, SRX340-lem
	SRX345	SRX345-dual-ac, SRX345-lem, SRX345, SRX345-dc
	SRX380	SRX380
	SRX 1500	SRX1500
	SRX4100	SRX4100
	SRX4200	SRX4200
	SRX4600	SRX4600
	SRX5000	SRX5000
SRX5400	SRX5400	
SRX5600	SRX5600	
SRX5800	SRX5800	
	vSRX Virtual Firewall, vSRX3	vSRX Virtual Firewall
Cisco	Cisco IOS XR NOTE: <ul style="list-style-type: none">• <i>Cisco Model Driven Telemetry (MDT) is not supported.</i>	-
Nokia	7250, 7750	

RELATED DOCUMENTATION

[Upgrade the Device Image | 156](#)

[View and Manage Device Configuration | 165](#)

[Zero-Touch Provisioning Overview | 132](#)

Add Devices

IN THIS SECTION

- [Discover Devices | 149](#)
- [Add New Devices | 153](#)

An administrator or a user with add device privileges can add devices to Paragon Automation.

Discover Devices

Ensure that you do the following before you start discovering devices:

- The device is configured with a management IP address that is reachable from Paragon Automation.
- You configure an additional IP address for the management Ethernet interface by including the `master-only` statement at the `[edit groups]` hierarchy level. You must use this additional IP address for onboarding the device. For more information, see [Management Ethernet Interfaces](#).
- A user with the privileges of an administrator is created and enabled on the device.
- For managed Nokia devices, ensure that:
 - Model-driven CLI is enabled by executing the following command in the device CLI:

```
configure system management-interface configuration-mode model-driven
```

- NETCONF is enabled on port 22, and you have NETCONF access by executing the following commands in the device CLI:

```
system management-interface netconf admin-state enable
system security user-params local-user user <user-name> access netconf
system management-interface netconf admin-state enable
system management-interface netconf port 22
```

- Sufficient inbound connections are allowed (base component establishes an SSH session to the device) by executing the following command:

```
system login-control ssh inbound-max-sessions <count>
```

- For Cisco devices, ensure that you enable NETCONF by committing the following configuration on the device.

```
xml agent
netconf agent tty
netconf-yang agent ssh
ssh server netconf vrf default
```

If the discovery of Cisco devices fail with the message connection reset by peer, set the SSH session rate limit to 600 on the Cisco devices.

```
host@device#configure
host@device(config)#ssh server rate-limit 600
host@device(config)#commit
```

To discover one or more devices:

1. Select **Configuration > Devices**.

The Devices page appears.

2. Click the **Add** icon (+).

The Add Devices page appears.

3. Click **Discover Devices**.

4. You can choose to either specify the details of devices manually, or import the details from a comma-separated values (CSV) file on your local computer.

Do one of the following:

- To specify the details manually, Click **Enter Manually**. This is the default option.

Complete the configuration according to the guidelines specified in [Table 36 on page 152](#).

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

- To import the device details by using a CSV file:
 - a. Click **Import From File**.
 - b. (Optional) Click **Download Sample CSV file** to download a sample CSV template (**DiscoverDeviceSample.csv**) and use it to specify the details of the devices to be discovered.
 - c. Save the file on your local system.
 - d. Click **Browse** to navigate to the folder and select the CSV file from your local file system.
 - e. Click **Open** to upload the CSV file.

The Targets and Credentials section on the Add Devices page displays the details that are automatically populated from the CSV file.

5. (Optional) Click the **Add** icon (+) to add the details of another set of devices to be discovered.

To delete one or more sets of devices and their credentials, select the corresponding Targets and Credentials section and click the **Delete** icon (trashcan).
6. Click **OK** to save your changes.

A confirmation message appears, indicating that a job is triggered to discover the devices
7. (Optional) You can click the job link in the confirmation message to view the status of the job on the jobs page (**Monitoring > Jobs**).

After the job completes successfully, the discovered devices are listed on the Devices page.

Table 36: Fields on the Add Devices Page - Discover Devices

Field	Description
Managed Status	<p>Click this toggle button to discover managed (default) or unmanaged devices.</p> <ul style="list-style-type: none"> • When you enable this option, Juniper Networks devices, Nokia devices and Cisco IOS XR devices in your network are discovered. • When you disable this option, device discovery is not triggered. However, all the details of the device is retrieved from the path computation element (PCE) and stored in the database. The Devices page displays the hostname and the IP address of the unmanaged devices.
Hostname/IP Targets	<p>Enter the hostname or IPv4 addresses of the devices to be discovered. You can specify a combination of individual IPv4 addresses, a subnet (for example, 10.0.0.1/24) or a range of IPv4 addresses (for example, 10.0.0.1 - 10.0.0.20).</p> <p>When you specify a subnet or a range of IPv4 addresses, all the devices that have the IPv4 addresses within that subnet or range are discovered.</p>
Add targets from topology to this list	<p>Click this link to add the details of unmanaged devices that are discovered by the path computation element in a network, to the database.</p>
Device Credentials	<p>Enter the username and password of the devices you are discovering.</p> <p>If the username and password for the devices that you are discovering are different for each device, use different Targets and Credentials section to provide the targets for discovering the devices.</p> <p>Specify targets and corresponding device credentials for discovering different devices by clicking the Add icon.</p>

Table 36: Fields on the Add Devices Page - Discover Devices (Continued)

Field	Description
Use Same Credentials for Managing the Device	<ul style="list-style-type: none"> Manage a device using RADIUS (or custom) credentials. Enter the credentials in the Device Credentials (Username and Password) fields and then enable the Use same credential for managing toggle button. When enabled, Paragon Automation uses the credentials you entered to connect to and manage the device. NOTE: To use RADIUS authentication on the device, you must configure information about the RADIUS servers on the network. For more information, see Radius Authentication. Manage a device using the Paragon Automation generated credentials. To use the Paragon Automation generated credentials for discovering and managing a device, disable the Use same credential for managing toggle button. When disabled, Paragon Automation generated credentials are used for all connection to the device and the Username and Password fields are prepopulated with the credentials generated by Paragon Automation.

Add New Devices

Before you add a new device, ensure that:

- You configure an additional IP address for the management Ethernet interface by including the master-only statement at the [edit groups] hierarchy level. You must use this additional IP address for onboarding the device. For more information, see [Management Ethernet Interfaces](#).

To add new devices:

- Click the **Add** icon (+).
The Add Devices page appears.
- On the Add Devices page, click **Add New Devices**.
- Specify the root password and the range of IP addresses for management connectivity, according to the guidelines specified in [Table 37 on page 154](#).
- You can choose to either specify the details of devices manually, or import the details from a comma-separated values (CSV) file.

Do one of the following:

- To specify the details manually, click **Enter Manually**. This is the default option.

Complete the configuration according to the guidelines specified in [Table 37 on page 154](#).

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

- To import the device details from a CSV file:
 - a. Click **Import From File**.
 - b. (Optional) Click **Download Sample CSV file** to download a sample CSV template (**DiscoverDeviceSample.csv**) and use it to specify the device details that you can import.
 - c. Save the file on your local system.
 - d. Click **Browse** to navigate to the folder and select the CSV file from your local file system.
 - e. Click **Open**.

The *Device Model* section on this page displays the details that are automatically populated from the CSV file.

5. (Optional) Click the **Add** icon (+) to add the details of another set of devices to be added.

To delete one or more sets of devices and their credentials, select the corresponding section and click the **Delete** icon.

6. Click **OK** to save your changes.

A confirmation message appears, indicating that a job is triggered to discover the devices.

7. (Optional) You can click the job link in the confirmation message to view the status of the job on the jobs page (**Monitoring > Jobs**).

After the job completes successfully, the devices are listed on the Devices page.

Table 37: Fields on the Add Devices Page-Add New Devices

Field	Description
-------	-------------

Management Connectivity

Table 37: Fields on the Add Devices Page-Add New Devices (*Continued*)

Field	Description
Root Password	<p>Click to enable or disable the toggle button.</p> <ul style="list-style-type: none"> When you enable this toggle button, the root password, that you set, is assigned to all the new devices being added. <p>Follow the Junos OS password rules for the device model being added to set the password; the password is usually 6 to 128 characters long, and contains mixed case letters combined with numbers and/or symbols.</p> <ul style="list-style-type: none"> When you disable, you must assign the password to each device separately.
IP/Subnet/Gateway	<p>Enter the range of addresses that can be configured on the devices being added. These IP addresses are added to a DHCP server and then assigned to the devices.</p>

Device Models

Click the **Add** icon (+) to add more device models for discovery.

Device Family	Select the device family that you want to add from the list; for example, Juniper-ACX.
Device Model	Select the device model that you want to add; for example, ACX2000.
Junos Image	Select the Junos OS image that the device must use. The default is Use Image on Device indicating that the device is added to Paragon Automation with the image already existing in it.
Device Serial Numbers	<p>Enter the serial number of one or more devices to be added.</p> <p>To add more than one serial number, enter the serial number and then press Enter.</p>
Root Password	When the common root password is disabled, enter the root password to be assigned to the device. Follow the Junos OS password rules for the device model being added to set the password; the password is usually 6 to 128 characters long including mixed case letters combined with numbers and/or symbols.

SEE ALSO

[View and Manage Device Configuration | 165](#)

[Zero-Touch Provisioning Overview | 132](#)

[Edit Devices | 168](#)

Upgrade the Device Image

For devices to support the latest features, you must upgrade the image running on the device to the latest available image. Refer to "[Image Upgrade Workflow](#)" on page 188.

To upgrade images on one or more devices:

1. Select the devices that you want to upgrade and click **More > Upgrade Devices**.

The Upgrade Devices page appears.

2. Do the following:

- To upgrade a different image on one or more devices, select the devices for which you want to choose an image to upgrade and click the **Edit** (pencil) icon.

The Selected Image field for each device lists the images that you can choose to upgrade.

- To upgrade the same image on two or more devices, select the devices and click **Bulk Select Image**. Select the image that you want to upgrade in the Select Image page that is displayed.

NOTE: While upgrading the image of a device, you cannot copy the image on to the device if the bandwidth on the device is lesser than 600Kbps.

3. For each device, select the image that you want to upgrade from the list.

4. In **Upgrade Devices**, click:

- **Run Now** to upgrade the image immediately.
- **Schedule Later** and select the date and time to schedule the upgrade.

5. Click **OK**.

A job is created to upgrade the image. You can view the progress of the job by clicking the job ID that appears on the top of the page or by navigating to the Jobs page (Monitoring > Jobs).

RELATED DOCUMENTATION

[Stage an Image | 193](#)

View Device Statistics and Inventory information

IN THIS SECTION

- [View Device Statistics | 157](#)
- [View Device Inventory | 159](#)

You can use the *Device-Name* page to view statistics, inventory and the configuration templates assigned to the device.

To access this page:

1. Select **Configuration > Devices**.

The Device page appears.

2. Click a device in the Device Name column of the Devices list.

The *Device-Name* appears displaying the Overview, Inventory, and Configuration Template tabs.

View Device Statistics

To view the device statistics, click the **Overview** tab on the *Device-Name* page to view information about the device. The overview tab displays the Chassis View and the widgets listed in [Table 39 on page 159](#).

You can view the following information on the Overview tab of the *<device-name>* page:

- The Chassis View, present at the top of the Overview tab, displays the device ports as icons. Hover over the ports icon to view the details, such as input and output bandwidth, port mode, negotiated speed, input and output errors, and admin status of the port.

For a Virtual Chassis, Chassis Cluster, or dual Routing Engine, the member devices are listed on the left side of the chassis view. Clicking the member devices displays the ports of the member device.

NOTE: When you click a device on the Devices page, you cannot view the chassis details if the menu bar, which is available on the left-side of the Paragon Automation GUI, is expanded. Minimize the menu bar to view the chassis details.

Chassis view is not supported for Nokia and Cisco-IOS XR devices.

At the top of the Chassis view, system meters are provided to display data about functioning of the device components such as CPU, memory, storage and fan. If data is not collected for the system meters (for example, due to issue with telemetry manager service), dummy data is displayed.

[Table 38 on page 158](#) lists the system meters present on the top right-corner of the Chassis View.

Table 38: Widgets on the Overview Tab

System Meter	Description
CPU	Displays the percentage of CPU utilized and that is idle in the device
Memory	Memory (in %) utilized in the device.
Storage	Storage space (in %) allocated to the logical partitions of the device (/ , /var, and /temp) NOTE: The data displayed for storage is dummy data for all devices.
Fan	Details of fans present in the device.
Temperature	Temperature of the FPCs in the device.

- The widgets display information related to port status, general device details, alarms, and resource utilization.

[Table 39 on page 159](#) describes the widgets on the Overview tab.

Table 39: Widgets on the Overview Tab

Widget	Description
Port Count by Status	<p>Graphical representation (Donut chart) of the port status.</p> <p>Hover over the chart to view the number and percentage of ports that are up and down.</p> <p>Click the More Details link (at the bottom, right-corner of the widget) to view the hardware and the interface details on the Inventory tab of the <i>Device-Name</i> page.</p>
Recent Alarms	<p>Lists the major, minor, and normal alarms generated recently on the device. Click the Show All link (at the top, right-corner of the widget) to view information about alarms generated during the past 1 hour, 8 hours, 1 day, 1 week, and 1 month.</p> <p>Click the View all alarms link to view the alarms on the Alarms page.</p>
Device Details	<p>Displays general details such as serial number, device platform, device family, OS version, management status, and configuration status of the device.</p>
Resource Utilization	<p>Graphical representation of memory and CPU utilized in the device for the selected time period (past 1 hour, 8 hours, 1 day, 1 week, and 1 month).</p> <p>NOTE: Resource utilization widget is not displayed for Cisco IOS XR devices.</p>

You can perform the following actions from the Overview tab:

- View only specific widgets—Click the **More Options** icon (vertical ellipses) on the right side of the page. Select the check boxes corresponding to the widgets that you want to view.

Only the widgets that you selected appear on the Overview tab.

- Refresh data on the widgets—To obtain the most current device data and statistics, click the Refresh icon (circular arrow) on the right side of the page.

The data on the widgets is refreshed.

View Device Inventory

To view inventory of a device, click the **Inventory** tab on the *<Device-Name>* page. A device inventory comprises hardware components (Chassis, FPC, PIC, and Interfaces) and software components (Licenses, Features, and Software).

Paragon Automation displays the inventory of a device containing data retrieved from the device during discovery and resynchronization operations. This inventory includes the number of available slots, power supplies, chassis cards, fans, part numbers, and so on. Sorting is disabled on the Chassis tab to preserve the natural slot order of the devices.

Paragon Automation displays the inventory information for primary and the member devices for dual Routing Engines, Chassis Clusters, and Virtual Chassis.

NOTE: For Juniper Networks and Cisco IOS XR devices, the inventory is synchronized whenever a change in the components is indicated in the device system log.

The inventory tab displays the following information:

- Chassis—View the hierarchical list of all the hardware components present on the chassis, and the associated physical interfaces and logical interfaces. [Table 40 on page 161](#) describes the fields displayed on the Chassis tab.
- Interfaces—View details of the interfaces present on the chassis, line card, FPC, and MICs. [Table 41 on page 162](#) describes the fields displayed on this tab.

You can perform the following functions on this tab:

- Change the administration status of an interface by selecting the interface and clicking **More > Change Admin Status**.

A message indicating the change in the admin status is initiated is displayed along with the job ID. Click the link job ID link to view the progress of the job on the Jobs page (Monitoring > Jobs). The admin status is changed after the job completes successfully.

NOTE: For Nokia devices, you cannot change the admin-status of a port from the **Inventory** tab.

- View the following details of an interface, as widgets, by clicking on the interface link:
 - General details, such as the port name, IP address, admin status, link status, and so on of the interface.
 - Bandwidth consumed by the interface for the previous 1 hour, 8 hours, 1 day, 1 week, and 1 month.
 - Errors that occurred on the interface during the previous 1 hour, 8 hours, 1 day, 1 week, and 1 month.

- Licenses—View details (such as the license version, validity type, and start and end dates) of the licenses installed on the device. [Table 42 on page 163](#) describes the fields on this tab.

NOTE: You cannot view the License information for Cisco IOS XR and Nokia devices.

- Features—View details (such as feature name, validity type, and licenses) of the features that are configured on the device. [Table 43 on page 164](#) describes the fields on this tab.

NOTE: You cannot view the feature information for Nokia devices.

- Software—View details (such as software name, version, and state) of the software installed on the device. [Table 44 on page 165](#) describes the fields on this tab.

NOTE: You cannot view software information for Nokia devices.

- On all the tabs, click the **Show or Hide Columns** icon and select the check boxes corresponding to the columns that you want to view.

If you want to revert to the default column selections, click **Reset Preferences**. The selections that you previously made are cleared, and the default columns are displayed.

[Table 40 on page 161](#) describes the fields on the Chassis tab.

Table 40: Fields on the Chassis Tab

Field	Description
Module	Name of the hardware component (chassis, FPC, MIC, midplane, and so on).
Type	Type of the hardware component—Chassis, line card, CPU, backplane, and so on.
Version	Version of the hardware component.
Part Number	Part number of the hardware component. Built-in indicates that the PIC is a part of the parent component and does not have a part number.

Table 40: Fields on the Chassis Tab (Continued)

Field	Description
Serial number	Serial number of the hardware component.
Physical Interfaces	Click the View link displayed in this column to view the physical interfaces associated with the hardware component. You are directed to the Interfaces tab. See Table 41 on page 162 for details.
Description	Brief description of the hardware component.

[Table 41 on page 162](#) describes the fields on the Interfaces tab.

Table 41: Fields on the Interfaces Tab

Field	Description
Interface Name	Name of the interface. You can click an interface name to view more information (such as additional details, bandwidth, and errors) on the widgets that appear. Click the Interface link present above the widgets to return back and view the list of all the interfaces on the device.
ifIndex	The SNMP object identifier.
MAC Address	MAC address of the interface.
Admin Status	Administration status of the interface—Up or Down.
Operational Status	Operational status of the interface—Up or Down.
MTU	Maximum transmission unit (MTU) size (in bytes) on the interface.
Speed	Interface speed in Mbps or Gbps.
Duplex Mode	Indicates if the duplex mode on the interface is full-duplex or half-duplex.

Table 41: Fields on the Interfaces Tab (Continued)

Field	Description
Link Type	Type of link provided by the interface. Example: Ethernet, Fast Ethernet
Line card	The line card number in which the interface is present.
Logical Interfaces	Number of logical interfaces associated with the physical interface. Click the number displayed in the column to view the logical interfaces and their details. The details are displayed on the Logical Interfaces page.

[Table 42 on page 163](#) describes the fields on the Licenses tab.

Table 42: Fields on the Licenses Tab

Field	Description
License Name	Name of the license.
Version	Version of the license.
State	State of the license: <ul style="list-style-type: none"> • Valid: The installed license key is valid. • Invalid: The installed license key is invalid. • Expired: The license key is expired.
Validity Type	Validity type of the license: <ul style="list-style-type: none"> • Databased: License expires on end date. • Permanent: License never expires. • Countdown: License expires when time remaining is zero • Trial: License expires at the end of the trial period.

Table 42: Fields on the Licenses Tab (Continued)

Field	Description
Start Date	Start date of the license.
End Date	End date of the license.
Features	Features that the license supports.

[Table 43 on page 164](#) describes the fields on the Features tab.

Table 43: Fields on the Features Tab

Field	Description
Feature Name	Features available on the device.
Description	Brief description of the feature.
Used Count	Number of licenses currently being used on the device. Usage is determined by the configuration on the device. If a feature license exists and that feature is configured, the license is considered used.
Installed Count	Number of licenses installed on the device for the particular feature.
Need Count	Number of licenses required for legal use of the feature. Usage is determined by the configuration on the device:
Validity Type	Validity type of the license. <ul style="list-style-type: none"> • Databased: License expires on end date. • Permanent: License never expires. • Countdown: License expires when time remaining is zero • Trial: License expires at the end of the trial period.

Table 43: Fields on the Features Tab (Continued)

Field	Description
Licenses	Number of licenses installed for the feature.

Table 44 on page 165 describes the fields on the Software tab.

Table 44: Fields on the Software Tab

Field	Description
Software Name	Name of the software running on the device.
Software Description	Brief description of the software.
Version	Software version.
State Type	The type of the component. The value is always set to Software Component.

SEE ALSO

[Zero-Touch Provisioning Overview | 132](#)

[Add Devices | 149](#)

[Deploy a Configuration Template to a Device | 293](#)

View and Manage Device Configuration

You can view and manage device configuration from the Configuration Template tab on the *Device-Name* page.

To view and manage a device configuration:

1. Select **Configuration > Devices** in the left navigation menu.

The Devices page is displayed.

2. Click the *Device-Name* link.

The *Device-Name* page appears.

3. Click Configuration Template tab to view the device configuration.

The Configuration Template tab displays the predefined and user-defined configuration templates assigned to the device. [Table 45 on page 166](#) describes the fields on this tab. See [Table 58 on page 279](#) to view the list of default configuration templates provided by Paragon Automation.

Here's what you can do on this tab:

- View Active Configuration—To view the configuration currently active on the device, click **View Active Configuration**. The Active Configuration for *Device-Name* page appears, displaying the configuration in the Set, Junos Native, and XML formats for Juniper Networks devices.

For devices from other vendors, you can view the active configuration in the Set and XML formats only. Click the close icon (X) or **OK** to close this page.

- Configure template parameters—To configure template parameters for the device, select a template from the list and click **Configure**. The Template Parameters page appears. Configure the parameters on the Configure tab and then, preview the configuration on the Preview tab. Click **OK** to save the configuration.

The Deployment Status changes to **Ready to Deploy**.

- Validate the configuration template—To validate the configuration template before deploying it on the device, select the configuration template, and click **Validate**. A job is created to validate the configuration template.
- Deploy the configuration template on the device—To deploy one or more configuration templates on the device, select the templates and click **Deploy**. A job is created to deploy the configuration. You can view the status of the jobs from the Jobs page (**Monitoring > Jobs**).

Table 45: Fields on the Configuration Templates Tab

Field	Description
Name	Name of the configuration template.

Table 45: Fields on the Configuration Templates Tab (*Continued*)

Field	Description
Deployment Status	<p>Deployment status of the configuration template.</p> <ul style="list-style-type: none"> • No Configuration—Indicates that either the configuration template is not applied to the device or the configuration template does not have any values assigned to the parameters. • Deployed—Indicates that the configuration template is deployed on the device • Ready to Deploy—Indicates that the configuration template is configured on the system and is ready to be pushed on to the device • Deployment-In-Progress—Indicates that the configuration template is currently getting deployed on the device. • Deployment Failed—Indicates that the configuration template push failed on the device (due to some errors or conflicting configuration)
Deployment Date	Date and time at which the configuration template was deployed on the device in the Month DD, YYYY hh:mm:ss AM/PM format.
Description	Description of the configuration template.
Validation	Indicates whether the validation of the configuration template was successful or not.

RELATED DOCUMENTATION

[Deploy a Configuration Template to a Device | 293](#)

[Zero-Touch Provisioning Overview | 132](#)

[Upgrade the Device Image | 156](#)

Edit Devices

You must be an administrator or a user with edit device privileges to edit a device. You edit the device parameters to update the Paragon Automation database with values of device parameters that are not captured during device discovery or synchronization.

To edit a device configuration in the database:

1. On the navigation menu, select **Configuration > Devices**.

The Devices page appears.

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

The Edit *Device-Name* page appears.

3. Edit the parameter values by referring to [Table 46 on page 168](#).

4. Click **OK**.

A message indicating that the device parameters are successfully updated in the Paragon Automation database is displayed, and you are returned to the Devices page.

Table 46: Editable Fields on the Edit *Device-Names* Page

Field	Description
Device Group	<p>The device group to which the device belongs. By default, all devices discovered in a topology are assigned to the controller group, A device must be assigned to the controller group if you want to manage label-switched paths (LSPs) on the device.</p> <p>NOTE: This field is not in use in this release, .</p>
<i>Authentication</i>	
Authentication Method	<p>Displays the authentication method—Credential-Based.</p> <p>If you have enabled the Use same credential for managing toggle button while adding a device, you can enter the RADIUS (or custom) credentials in the Username and Password fields.</p> <p>If you have disabled the Use same credential for managing toggle botton while adding a device, then the credentials generated by Paragon Automation are automatically populated in the Username and Password fields.</p> <p>The change in credentials does not affect the active session, and Paragon Automation must establish a new session to connect to the device for the changes to take effect.</p>

Table 46: Editable Fields on the Edit *Device-Names* Page (Continued)

Field	Description
Username	<ul style="list-style-type: none"> Enter the username to access the device with RADIUS (or custom) credentials. Enter the username to access the device with credentials you provided. <p>The credentials you provided in the Username field will replace the credentials prepopulated by Paragon Automation. To not lose access to the device, the credentials you provided must be configured on the device.</p>
Password	<ul style="list-style-type: none"> Enter the password to access the device with RADIUS (or custom) credentials. Enter the password to access the device with credentials you provided. <p>The credentials you provided in the Password field will replace the credentials prepopulated by Paragon Automation. To not lose access to the device, the credentials you provided must be configured on the device.</p>
<i>Protocols:SSH</i>	
Timeout	<p>Number of seconds before a connection request to the device times out. The default is 300 seconds. Use the up arrow to increment and the down arrow to decrement this value or type a different value in the field.</p> <p>A value of 0 seconds allows a persistent connection (that is, the connection is available forever).</p>
Maximum Retry Count	Number of retries for establishing a connection with the device.
<i>Protocols:SNMP</i>	
Version	Select the SNMP version to be used—SNMP v2 or SNMP v3. The default is SNMPv2.
Get Community	Enter the SNMP get community string as configured on the device. The default is <i>public</i> .

Table 46: Editable Fields on the Edit *Device-Names* Page (Continued)

Field	Description
Port	Port number on which the SNMP requests are to be sent. The default is 161.
Timeout	Enter the number of seconds after which an SNMP connection or request times out. The default is 3 seconds.
Retry Count	Number of times that an SNMP connection can be attempted. The default count is 3.
<i>Protocols:PCEP</i>	
Version	Select a Path Computation Element Protocol (PCEP) version to use from the list. <ul style="list-style-type: none"> • Non-RFC Select this version to run in non-RFC 8231/8281 compliance mode. This is the default option. • RFC Compliant Select this version to run in RFC 8231/8281 compliance mode. This is supported in Junos OS 19.x and later (Junos OS releases that are RFC 8231/8281 compliant) releases. • Third-party path computation client (PCC) Select this version for any non-Juniper device that does not support association object.
IP Address	Enter the PCEP IP address used by the device to connect to Paragon Automation for managing LSPs. The PCEP IP address is usually the management IP address of the device.
<i>Protocols:PRPD</i>	
Enabled	Click to enable or disable (default) the Programmable Routing Protocol process (PRPD) on the device.

Table 46: Editable Fields on the Edit *Device-Names* Page (Continued)

Field	Description
Enable SSL	Click to enable or disable (default) use of SSL by PRPD.
IP	IP address for PRPD on the device. The default is the device's loopback address. If you do not enter a value, the device's loopback address is used.
Port	Port on the device to establish a PRPD session. The default is 50051.
<i>Protocols:Netconf</i>	
Enabled	Click to enable or disable NETCONF on the device.
Bulk Commit	Click to enable or disable bulk commit on the device. If you enable this option, you can provision multiple LSPs in a single commit instead of using multiple commits. This improves provisioning efficiency. This option should be enabled for Point-to-Multipoint-Traffic Engineering (P2MP-TE) when you use P2MP on Juniper Networks devices.
Retry Count	Enter the number of times that a connection can be attempted on the device. The default is 3. A value of 0 means an unlimited number of retries; connection attempts never stop.
iAgent/Netconf Port	Enter the port number for NETCONF on the device. This port should not be used for any other service. The default port number is 830 for Juniper Networks devices and 22 for other devices.
<i>Device ID Details</i>	

Table 46: Editable Fields on the Edit *Device-Names* Page (Continued)

Field	Description
System ID to use for JTI	<p>Unique system identifier required for Junos Telemetry Interface (JTI) native sensors. Junos OS devices use the following format: <i>host_name.jti_ip_address</i> for System ID where, <i>host_name</i> is the host name of the device and the <i>jti_ip_address</i> is the (local-address statement) that is configured for the export profile in Junos OS.</p> <p>When a device has dual routing engines (REs), it might send different system IDs depending on which RE is primary. You can use a regular expression to match both system IDs.</p>
Management IP	Enter the management IP address of the device.
Flow/IFA Source IPs	<p>Enter one or more IP addresses (separated by commas) that the device uses to send NetFlow data to Paragon Automation.</p> <p>The IP address or addresses are used to send probe packets for flow monitoring using Inband Flow Analyzer (IFA).</p> <p>If there are more than one IP address, separate them with a comma.</p>
Syslog Source IPs	Enter the list of IP addresses (separated by commas) that the device uses to send system log messages to the Paragon Insights component of Paragon Automation. For example, 10.10.10.23, 192.168.10.100.
Syslog Host Names	Enter the list of hostnames (separated by commas) for sending system log messages to the Paragon Insights component of Paragon Automation.
SNMP Source IPs	Enter the list of IP addresses (separated by commas) for sending SNMP messages to the Paragon Insights component of Paragon Automation.
<i>Advanced</i>	
Initial Sync	<p>Click this toggle button to enable or disable (default) the device from sending a continuous stream of data to Paragon Insights until the device and Paragon Insights are synchronized.</p> <p>This option is disabled by default to reduce the processing load on the device.</p>

Table 46: Editable Fields on the Edit *Device-Names* Page (Continued)

Field	Description
gNMI Support	Click this toggle button to enable or disable (default) the gRPC network management interface (gNMI) on the device.
gNMI Encoding	If you enable gNMI support, select the gNMI encoding to be used from: <ul style="list-style-type: none"> • protobuf (default) • json • json_ietf
Open Config Port	Enter the port on which the gRPC connection needs to be established for OpenConfig telemetry.
Syslog Time Zone	Specify the time zone in the format +/- hh:mm (with reference to GMT) for parsing the time stamp in the system logs. The system log time zone is usually the time zone in which the device is located.

RELATED DOCUMENTATION

[Zero-Touch Provisioning Overview | 132](#)

[Add Devices | 149](#)

[Upgrade the Device Image | 156](#)

Delete Devices

You must be an administrator or a user with privileges to delete a device.

Before you delete a device, ensure that the device is not a part of any device group. If the device is a part of a device group, unassign the device from the device group.

To delete a device from Paragon Automation:

1. Select **Configuration > Devices** on the navigation menu.

2. Select one or more devices that you want to delete and click the **Delete** icon (trashcan).

A confirmation message appears asking you to confirm the delete.

3. Click **Yes**.

A job is created to delete the device and the job ID appears on the top of the Devices page. You can view the progress of the job in the Jobs page (Monitoring > Jobs). The device is removed from Paragon Automation and no longer listed on the Devices page after the job completes successfully.

Sometimes, a device might be referenced in other components of Paragon Automation. In such a case, you cannot delete a device. If you still have to delete the device, you must force delete the device.

To force delete a device, follow the steps for a normal delete. In the confirmation message that appears, click the **Force Delete** option and click **Yes**. A job is created to delete and the device is deleted from the Paragon Automation database.

RELATED DOCUMENTATION

[Zero-Touch Provisioning Overview | 132](#)

[Add Devices | 149](#)

[Upgrade the Device Image | 156](#)

Device Groups

IN THIS CHAPTER

- [About the Device Groups Page | 175](#)
- [Add a Device Group | 177](#)
- [Edit a Device Group | 183](#)
- [Filter a Device Group | 183](#)
- [Delete a Device Group | 184](#)
- [Commit or Roll Back Configuration Changes in Paragon Insights | 185](#)

About the Device Groups Page

IN THIS SECTION

- [Tasks You Can Perform | 176](#)

You must add a device to one or more device groups before you can apply rules and playbooks to a device. When you add a device group, you can perform the following tasks:

- Select devices to be included in the device group
- Set retention policy for time series database
- Set port numbers for Native GPB, Syslog, SNMP Notification, and NetFlow sensors
- Select report profiles for a device group
- Configure SNMP ingest
- Select ingest frequency profile

- Configure logging for supported sensors
- Apply tagging and summarization (data rollup and raw data) profiles
- Configure SSH or TLS authentication
- Select notification groups to publish sensor and field data

To access device group page in Paragon Automation Platform GUI, click **Configuration > Device Groups**.

[Table 47 on page 176](#) describes the fields in the Device Groups page.

Table 47: Fields in Device Groups Page

Fields	Description
Device Group Name	Displays the name of the device group.
Devices	Displays the name of a device and '+' followed by a number. If you hover over the number, it displays the names of other devices added to the device group.
Playbooks	Displays the name of a playbook and '+' followed by a number. If you hover over the number, it displays the names of other playbooks deployed in the device group.
Status	Displays the deployment status of the device group. If a device group is not deployed, it shows Saved.
Logging	<p>Displays the first letter of the log severity such as Debug, Error, Warn or Info along with the log level.</p> <p>The log level for the device group shows Global if you configure device group to collect logs for every service running on a device group.</p> <p>If you configure service specific logs, the column displays Others.</p>
Notifications	Displays the number of alerts generated by each device group. If you hover over the number, it shows the breakdown of alerts by severity.

Tasks You Can Perform

You can perform the following tasks in this page:

- Add a device group. See ["Add a Device Group" on page 177](#).
- Edit a device group. See ["Edit a Device Group" on page 183](#).
- Delete a device group. See ["Delete a Device Group" on page 184](#).
- Filter and Search a device group. See ["Filter a Device Group" on page 183](#).
- Export configuration details of device groups.

To export configuration details of all device groups:

1. Go to **Configuration > Device Groups**.

You are taken to the Device Group Configuration page.

2. Click on the **Export** button and select *Export as CSV* from the menu.

In the pop up window, click **Open** to view the Excel file or **Save As** to save the Excel file to any location in your system.

RELATED DOCUMENTATION

| [Enable Alert Notifications for Device Groups and Network Groups | 627](#)

Add a Device Group

You must add a device to one or more device groups before you can apply rules and playbooks to a device.

To add a device group:

1. Click **Configuration > Device Groups**.

The Device Group Configuration page appears.

2. Click the add (+) icon.

The Add Device Group page appears.

3. Configure the device group fields as per [Table 48 on page 178](#).

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

4. Do one of the following:

- Click **Save** to save the configuration. Selecting this option does not deploy the new device group configuration.

NOTE: You can click the Health Configuration Deployment Status icon on the top-right corner to commit, deploy, or rollback the configuration settings. For more information, see ["Commit or Roll Back Configuration Changes in Paragon Insights"](#) on page 185.

A confirmation message appears stating that the device group is successfully created. The new device group is added to the Device Group page.

- Click **Save and Deploy** to save and immediately deploy the device group configuration.

A confirmation message appears that device group is successfully created and deployed.

Table 48: Fields on Add Device Group Page

Field	Description
General	
Name	Enter a name for the device group.
Description	Enter a description for the device group.
Devices	Select one or more devices to add to the device group.
Advanced	
Flow Ingest Deploy Nodes	Select one or more nodes from the existing nodes list to specify which nodes must receive the NetFlow traffic if your installation is a multi-node installation using Kubernetes.
Ingest Frequency Profiles	Select one or more profiles from the list that should override the rules or sensor frequency settings.
Reports	Select one or more health report profiles from the list to generate reports for the device group. Reports include alarm statistics, device health data, and device-specific information (such as hardware and software specifications).

Table 48: Fields on Add Device Group Page (*Continued*)

Field	Description
Retention Policy	<p>Select a retention policy from the list for time series data used for root cause analysis (RCA) in this device group. By default, the retention policy is 7 days. For more information, see "Configure a Retention Policy" on page 632 and "Time Series Database (TSDB) Overview" on page 635.</p>
<p>Logging</p> <p>You can collect logs of different severity level for a device group. Use the following fields to configure which log levels to collect.</p>	
Global Setting	<p>Select one of the following levels for log messages that you want to collect for the device group:</p> <ul style="list-style-type: none"> • None—Disables logging functionality • Error • Debug • Warn—Conditions that warrant monitoring • Info—Any events or non-error conditions of interest <p>NOTE: By default, the logging level is set to Error.</p>

Table 48: Fields on Add Device Group Page (Continued)

Field	Description
Service Logging Overrides	<p>Select the log level for any of the following services that you want to configure differently from the Global Log Level setting:</p> <ul style="list-style-type: none"> • Open Config • iAgent • Native GPB • Netflow • SNMP • SNMP Notification • Syslog • Reports Generation • Non-Sensor rules • Trigger Evaluation • ML Model Builder <p>NOTE: The log level that you select here for a specific service takes precedence over the Global Log Level setting.</p>
Notifications	<ul style="list-style-type: none"> • You can use the Notifications feature to organize, track, and manage KPI event alarm notifications received from devices in the device group. • To receive alarm notifications for KPI events that have occurred on your devices, you must first configure the notification delivery method for each KPI event severity level (Major, Minor, and Normal). <p>Select the delivery method from the respective lists.</p>
Ports	

Table 48: Fields on Add Device Group Page (Continued)

Field	Description
Native Ports	(Native GPB sensors only) Enter the port numbers on which the Junos Telemetry Interface (JTI) native protocol buffer connections are established as comma separated values.
Flow Ports	(NetFlow sensors only) Enter the port numbers on which the NetFlow sensor data is received as comma separated values. The port numbers must be unique.
sFlow Ports	(sFlow sensors only) Enter the port numbers on which the sFlow sensor data is received as comma separated values. The port numbers must be unique.
Syslog Ports	Enter the UDP ports on which syslog messages are received as comma separated values.
SNMP Notification Ports	Enter the SNMP Notification port numbers as comma separated values.
Outbound SSH Ports	Enter TCP port number(s) for NETCONF outbound SSH connections as comma separated values.
Summarization	<p>To improve the performance and disk space utilization of the time series database, you can configure data summarization methods to summarize the raw data collected.</p> <p>Use the following fields to configure data summarization:</p> <ul style="list-style-type: none"> • Profiles—Select the data summarization profiles which you want to apply to the ingest data. To edit or view details about saved data summarization profiles, see Summarization Profiles page under the Configuration menu. • Time Span—Enter the time span (in minutes) for which you want to group the data points for data summarization. <p>For more information, see "About the Data Summarization Page" on page 647.</p>
Data Rollup Summarization	Select the rollup summarization policy for the device group.

Table 48: Fields on Add Device Group Page (Continued)

Field	Description
Publish	<p>You can configure the PCE to publish sensor and field data to AMQP/Kafka servers for a specific device group:</p> <ul style="list-style-type: none"> Notification Groups—Select the notification groups for data publishing. <p>NOTE: Only kafka-publish and amqp-publish are currently supported.</p> <ul style="list-style-type: none"> Rules—Select the rule topic and rule name pairs that contain the field data you want to publish. Sensor—Select the sensor paths or YAML tables that contain the sensor data you want to publish. No sensor data is published by default.
Tagging	<p>Select one or more tagging profiles from the existing profiles list. Tagging makes use of profiles to set conditions, define new fields and keys, and insert values into those fields after creation. For more information, see "Paragon Insights Tagging Overview" on page 570.</p>
Root Cause Analysis	<p>Mode is enabled by default. Disable Mode if you do not want the device group to be a part of resource discovery and dependency process.</p> <p>Exclude Resources field allows you to select device resources that must be excluded from resource and dependency formation.</p>
IFA Deploy Nodes	<p>Enter IP address of nodes where In-band Flow Analyzer (IFA) ingest must be deployed in a multinode Paragon Insights installation.</p>
IFA Ports	<p>Enter a UDF port number from 1 to 65535.</p> <p>The port receives IFA sensor data in Paragon Automation.</p>

RELATED DOCUMENTATION

| [About the Device Groups Page](#) | 175

Edit a Device Group

To edit a device group:

1. Click **Configuration > Device Groups**.

The Device Groups page appears.

2. Select the device group that you want to edit.

3. Click the Edit (pencil) icon.

The Edit Device Group page appears.

4. Modify the device group attributes as per "[Add a Device Group](#)" on page 177.

NOTE: You cannot edit the Device Group Name field.

5. Do one the following:

- Click **Save** to save the configuration. Selecting this option does not deploy the updated device group configuration.

A confirmation message appears stating that the device group is successfully edited.

NOTE: You can click the Health Configuration Deployment Status icon on the top-right corner to commit, deploy, or rollback the configuration settings. For more information, see "[Commit or Roll Back Configuration Changes in Paragon Insights](#)" on page 185.

- Click **Save and Deploy** to save and immediately deploy the updated device group configuration.

RELATED DOCUMENTATION

[About the Device Groups Page](#) | 175

[Add a Device Group](#) | 177

Filter a Device Group

In the Device Group Configuration page, you can filter and search for a device group. You can also search for a device group by clicking on the search button (magnifying glass icon) and entering the name of a device group. The page displays the device group you entered in the search field.

To apply filter:

1. Go to **Configuration > Device Group configuration**.

You are taken to the Device Group Configuration page.

2. Click on the filter button (funnel icon) and click on the **Add Filter** button.

The Add Criteria window pops up.

3. Select a device, a device group, a playbook and so on from the drop-down menu of the Field box.

4. Select a condition from *Includes*, *=*, and *!=* from the drop-down menu of Condition box.

5. Enter the name of a device, device group, status (Deployed or Saved) in the Value box.

The results of your search is displayed on the Device Group Configuration page.

6. (Optional) Click on the **+** button at the side of the existing filter to add new filters.

7. (Optional) Click on the **Save** button to save your filter.

In the Save Filter window, enter a name for your filter and enable *Set as Default* if you want the filter to be marked as default.

You can find the saved and default filters when you click on the filter button. Click on the **x** beside the filter name to remove the filter.

RELATED DOCUMENTATION

[Add a Device Group | 177](#)

[Delete a Device Group | 184](#)

Delete a Device Group

To delete a device group:

1. Click **Configuration > Device Groups**.

The Device Group Configuration page is displayed.

2. Select the device group you want to delete.

3. Click the delete (trash can) icon.

A confirmation message appears.

4. Do one of the following:

- Click **Delete** to delete the selected configuration.

A confirmation message appears stating that the device group is successfully deleted. The selected device group is deleted from the Device Groups page.

- Click **Delete and Deploy** to delete and immediately deploy the device group configuration. This option will remove all the running instances associated with the device group.

A confirmation message appears stating that device group is successfully deleted.

RELATED DOCUMENTATION

| [Edit a Device Group](#) | 183

Commit or Roll Back Configuration Changes in Paragon Insights

In Paragon Insights, when you make changes to the configuration, you can perform the following operations:

- Save
- Save and Deploy
- Delete
- Delete and Deploy

These operations and their effect on the Paragon Insights ingest services are explained in [Table 49 on page 185](#).

For changes that were not already deployed, you can either commit or roll back the changes, which you can do using the Health Configuration Deployment Status page; the procedure is provided below [Table 49 on page 185](#).

Table 49: Operations After Changing the Configuration in Paragon Insights

Operation	Explanation	Effect on Ingest Services
Save	Saves the changes to the database, but doesn't apply the changes to the ingest services.	No effect until the changes are committed.
Delete	Deletes the changes from the database, but doesn't apply the changes to the ingest services.	

Table 49: Operations After Changing the Configuration in Paragon Insights (Continued)

Operation	Explanation	Effect on Ingest Services
Save and Deploy	Saves the changes to the database and applies the changes to the ingest services.	When you commit a configuration, the changes are accepted after validation and acknowledged immediately. A background job tracks these changes to ingest services.
Delete and Deploy	Deletes the changes from the database and applies the changes to the ingest services.	The background job runs periodically and applies the changes to ingest services. Only after the background job applies these changes, does the ingest service process data based on the changes configured.

NOTE: Users might perform the Save or Delete operations if they have multiple configuration steps and want to stack the configuration steps, and commit the changes later at one go.

To commit or roll back configuration changes in Paragon Insights:

1. On the Paragon Automation banner, click the deployment status icon.

The Health Configuration Deployment Status page appears displaying:

- The status of the last deployment.
- The list of device groups and network groups for which pending changes are not yet committed.
- The list of device groups and network groups for which pending deletions are not yet committed.

2. You can do one of the following:

- Commit any changes that are pending deployment:

a. Click **Commit**.

Paragon Insights triggers the commit operation and the configuration changes are applied to the ingest services. Depending on the number of configuration changes and number of device groups and network groups affected, applying the configuration changes to the ingest services might take between a few seconds or a few minutes to complete.

After the operation completes, a confirmation message appears. Paragon Insights starts running the associated playbooks and starts ingesting the data.

b. Go to Step 3.

- Roll back any changes that are pending deployment:

- a. Click **Rollback** to roll back any changes that are pending deployment. (The roll back operation discards the pending configuration changes that were previously saved in the database.)

Paragon Insights triggers the roll back operation and the configuration changes are rolled back almost immediately. After the rollback is complete, a confirmation message is displayed.

- b. Go to Step 3.

NOTE: After the roll back or commit operation completes, the status of the last deployment in the Health Configuration Deployment Status page displays the last operation that was completed.

3. Click **Close** to exit the page.

You are returned to the previous page from which you accessed the Health Configuration Deployment Status page.

RELATED DOCUMENTATION

[Add a Device Group | 177](#)

[Add a Network Group | 266](#)

Device Images

IN THIS CHAPTER

- [Image Upgrade Workflow | 188](#)
- [About the Images Page | 190](#)
- [Upload an Image | 192](#)
- [Stage an Image | 193](#)
- [Deploy an Image | 194](#)
- [Delete Images | 196](#)

Image Upgrade Workflow

The following is the image upgrade workflow in Paragon Automation:

1. Upload an image to Paragon Automation; see ["Upload an Image" on page 192](#).
2. Stage image on the device; see ["Stage an Image" on page 193](#).

Paragon Automation validates if the complete software is copied on to the device by using the checksum of the image. The checksum of the image in Paragon Automation is verified with the checksum of the image in the device.

If checksum of the image copied on to the device does not match with the checksum of the image in Paragon Automation, the image copied on to the device is deleted and the image is copied again. If the checksum does not match again, the stage task fails.

3. Deploy the image; see ["Deploy an Image" on page 194](#).

During deployment the following tasks are performed on the device:

- Validate the image copied on to the device—Paragon Automation validates if the complete software is copied on to the device by using the checksum of the image. The checksum of the image in Paragon Automation is verified with the checksum of the image in the device.

If checksum of the image copied on to the device does not match with the checksum of the device in Paragon Automation, the image copied on to the device is deleted and the image is copied again. If the checksum does not match again, the deploy job fails.

- Upgrade Image on the device—Paragon Automation upgrades devices as follows:
 - Single Chassis/Standalone devices: Normal upgrade, where the device stops forwarding traffic during the upgrade process.
 - Virtual Chassis (EX Series and QFX Series devices): Nonstop software upgrade (NSSU), where the image is copied on to the primary device and the primary device copies the image on to the member devices. The image is upgraded on the member devices one-by-one. The image on the primary device is upgraded last.

In this upgrade, there can be a minimal disruption to the traffic traversing through the device without affecting the control plane functioning.

- Chassis Clusters (SRX Series Firewalls): Both upgrade or downgrade is possible by using In-Band Cluster Upgrade (ICU) or In-Service Software Upgrade (ISSU);

See [Upgrading Devices in a Chassis Cluster Using ICU](#) for details about ICU and [Understanding In-Service Software Upgrade \(ISSU\)](#) for details about ISSU.

- vSRX Virtual Firewall Cluster: The primary and the backup Routing Engines are upgraded one by one. The secondary Routing Engine is upgraded first and switched over to the primary role. The former primary Routing Engine is then upgraded.

Traffic disruption is minimal during the upgrade.

- QFX and MX Series devices with dual Routing Engines: The primary and the backup routing engines are upgraded by Paragon Automation. Traffic is disrupted during the upgrade.
- Reboot the device—The devices are automatically rebooted after the image is upgraded.

Paragon Automation resynchronizes with the device after the device connects back.

NOTE: Paragon Automation does not support upgrade of Cisco-IOS XR devices. You must upgrade images on Cisco-IOS-XR devices manually.

RELATED DOCUMENTATION

[Upgrade the Device Image | 156](#)

[Deploy an Image | 194](#)

About the Images Page

IN THIS SECTION

- [Tasks You Can Perform | 190](#)
- [Field Descriptions | 191](#)

Paragon Automation helps you to manage (add, stage, deploy, and delete) the entire lifecycle of images of all managed network devices. A device image is a software installation package that you use to upgrade or downgrade the operating system running on a network device.

Paragon Automation can manage Junos OS images running on the following devices:

- ACX Series
- EX Series (both single and Virtual Chassis; mixed-mode Virtual Chassis supported)
- MX Series
- PTX Series
- QFX Series (both single and virtual chassis)
- SRX Series (both standalone and chassis clusters)
- vSRX

To access this page, click **Configuration > Device Images**.

Tasks You Can Perform

You can perform the following tasks from this page:

- Upload an image; see ["Upload an Image" on page 192](#).
- Stage an image; see ["Stage an Image" on page 193](#).
- Deploy an image; see ["Deploy an Image" on page 194](#).
- Delete images; see ["Delete Images" on page 196](#).
- Filter Entries—Filter the table entries by adding new filtering criteria.

Hover over the **Filter** (funnel) icon and select **Add Filter**. On the **Add Criteria** page, select the filtering criteria from the **Field** and **Condition** list, and enter the text to be compared in the **Value** field. Click **Add**.

The filtered table entries are listed and the filter criteria name is displayed above the table column names.

You can click the cross (X) icon (next to the filter name) to remove the filtering criteria.

NOTE:

- You can add multiple filtering criteria. Once you add the multiple filtering criteria, you can select the **And** condition to display the entries matching all the filtering criteria or select the **Or** condition to display the entries matching any one of the filtering criteria.
- Quick filter: Save the filtering criteria as quick filters. Once you have added all the filtering criteria, you can save a particular criteria or multiple criteria for future use by clicking **Save**.

On the **Save Filter** window, 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 a specific column in the table.

Hover over the **More Options** (vertical ellipsis) > **Show/Hide Columns** and select the *Column-Name* check box of the columns you want to display in the table.

- Reset Preference—Reset the displayed columns to the default set of columns for each tab in the table.

Hover over the **More Options** (vertical ellipsis) and select **Reset Preference**.

- Sort Entries—Click the column name to highlight the up and down arrows next to the column name. Sort the table entries in ascending or descending order of that column by clicking the up or down arrow respectively.

Field Descriptions

[Table 50 on page 192](#) displays the fields on the Images page.

Table 50: Fields on the Images Page

Field	Description
Image Name	The name of the image file.
Version	The version number of the image. For example: 20.4R1.12
Vendor	The vendor of the image. For example: Cisco Systems, Juniper Networks.
Family	The device family to which the image belongs. For example: Juniper-SRX, Juniper-EX
Supported Platform	Device models on which the image can be deployed. Only one device model (for example, MX40) is listed in this column. A + <Integer>, where, the integer indicates the number of additional device models supported is displayed next to the device model. Click the + <Integer> to view the list of all the other device models on which the image can be deployed.
Size	The size of the image file in MB or GB.
Uploaded By	The user who uploaded the image file.

RELATED DOCUMENTATION

[Devices Overview | 130](#)

[Upgrade the Device Image | 156](#)

Upload an Image

Before you begin, ensure that:

- You are able to access the image.
- You have the permission to upload the device image.

You upload images of devices to Paragon Automation so that you can manage the lifecycle of the image on a device. When you need to upgrade or downgrade the image running on a device, you can stage and deploy the required image on the device by using Paragon Automation.

To upload an image:

1. Click **Configuration > Device Images** in the left navigation menu.

The Images page appears.

2. Click the **Add (+)** icon.

The Upload Image page appears.

3. Do one of the following:

- To upload the image from your local system, click **Browse** under Upload from computer, and select the image file saved on your computer in the explorer that opens.

NOTE: When you try to upload a large image from your local system, the image upload might fail if there are firewall or issues in connecting with the network. In such scenarios, we recommend that you use the option to upload an image from the external server.

- To upload the image from an external server, enter the URL in the **Image URL** field; for example, you can provide the URL from the Juniper Networks Software Download page to download an image.

4. Click **OK**.

The image is copied to Paragon Automation server and listed on the Images page.

RELATED DOCUMENTATION

| [Upgrade the Device Image](#) | 156

Stage an Image

You need to stage or copy an image on to a device before upgrading the software running on the device. The Stage option is useful if you are using a low-bandwidth connection. On low bandwidth connections, manually staging an image prior to deploying the image helps prevent the image deployment from timing

out because of a slow connection. On high-bandwidth connections, you can choose to stage the image along with the image deployment.

When you stage an image, the checksum of the image in Paragon Automation is verified with the checksum of the image in the device. If checksum of the image on the device does not match with the checksum in Paragon Automation, the image copied on to the device is deleted and the image is copied again.

To stage an image, you must be an administrator or a user with the privilege to stage an image.

To stage an image:

1. Click **Configuration > Device Images** in the left navigation menu.

The images page appears.

2. Select the image that you want to stage and click the **Stage** button. You can stage an image onto a single device or multiple devices at the same time.

The Stage Image page appears.

3. Under Select Devices, select one or more devices onto which the image needs to be staged.

4. In the **Stage Image** field, click:

- **Run Now** to stage the image immediately.
- **Schedule Later** to stage the image later and specify the date and time when you want the image to be staged.

5. Click **OK**.

If you selected Run Now, a job is initiated immediately to stage the image. If you selected Schedule Later, a job is initiated at the scheduled date and time to stage the image. You can monitor the progress of the job in the Jobs page (Monitor > Jobs).

RELATED DOCUMENTATION

[Upgrade the Device Image | 156](#)

[Image Upgrade Workflow | 188](#)

Deploy an Image

To deploy an image, you should be an administrator or a user with the privilege to deploy images on devices. You can deploy an image on a single device or multiple devices at the same time.

NOTE:

- Upgrade of Junos Continuity software packages (JAM packages) is not supported.
- When you deploy an image on a device, the device goes into the maintenance state. In the maintenance state:
 - Other actions that impact the device such as, rebooting the device or deploying configuration templates is not possible.
 - Traffic flowing through an SRX Chassis Cluster or an EX Series or QFX Series Virtual Chassis is not disrupted.
 - Traffic flowing through a standalone device is disrupted.

NOTE: You can also upgrade images from the Devices page; see ["Upgrade the Device Image" on page 156](#).

To deploy an image onto a device:

1. Click **Configuration > Device Images** on the left navigation menu.

The images page appears.

2. Select the device image to be deployed on the device and click the **Deploy** button.

The Deploy Images page appears. In the Deploy Images page, you can view whether the image is staged on a device or not. If the image is not staged already, the image is copied on to the device and then deployed and this results in a longer deployment time.

3. Under Select Devices, select one or more devices onto which the device image needs to be deployed.

4. In the **Deploy Image** field, select a time to run the deployment:

- Click **Run Now** to deploy the image immediately.
- Click **Schedule Later** to deploy the image later and specify the date and time when you want the image to be deployed.

5. Click **OK**.

If you selected Run Now, a job is initiated immediately to deploy the image. If you selected Schedule Later, a job is initiated at the scheduled date and time to deploy the image. You can monitor the progress of the job in the Jobs page (Monitor > Jobs).

RELATED DOCUMENTATION

[Upgrade the Device Image | 156](#)

[Image Upgrade Workflow | 188](#)

Delete Images

You can delete one or more device images from the Images page when you no longer need to manage them.

To delete an image, you should be an administrator or a user with the privilege to delete an image. If you delete an image while the image is being staged or deployed, the job initiated to stage or deploy the image fails.

To delete an image:

1. Click **Configuration > Device Images** on the left navigation menu.

The images page appears.

2. Select one or more images that you want to delete and click the **Delete** icon (trashcan).

A confirmation message appears.

3. Click **Yes** to delete the images.

The selected images are deleted and the images are no longer listed on the Images page.

RELATED DOCUMENTATION

[Upload an Image | 192](#)

[Upgrade the Device Image | 156](#)

[Image Upgrade Workflow | 188](#)

Network

IN THIS CHAPTER

- Assign Names to Admin Group Bits | 197
- Modify Pathfinder Settings From the Pathfinder CLI | 198
- Modify Pathfinder Settings From the GUI | 206
- Disaster Recovery Overview | 252
- Network Slicing Overview | 256
- Add a Test Agent for Network Slices | 259
- Configure LSP Routing in a Network Slice by Using a Path Computation Profile | 261

Assign Names to Admin Group Bits

Admin groups, also known as link coloring or resource class assignment, are manually assigned attributes that describe the “color” of links, such that links with the same color conceptually belong to the same class. The path computation element (PCE) automatically discovers admin group bits associated with a link or label-switched path (LSP).

You can use the Admin Group page to assign a meaningful name to the admin group bits (0 through 31).

For example, you might want to map an admin group to a network region, such as San Francisco or New York, where each region is given its own bit name. Or, you can map each admin group to a color so that you can easily differentiate the different traffic routes in the topology map.

NOTE: You can access this page only if you have the permission to perform tasks (create, read, update, delete) in the Network page (**Configuration > Network**).

To assign names to admin group bits:

1. Select **Configuration > Network Settings > Admin Group**.

The Admin Group page appears.

2. Specify a meaningful name for one or more corresponding admin group bits (0 through 31).
If you don't specify a name, the default naming (bit0, bit1, and so on) is used.
3. Click **Save Changes**.
A confirmation message appears on the top of the page indicating that the changes have been saved successfully.

In the Link tab of the network information table, the admin group bit names appear in the Admin Group AZ and Admin Group ZA columns. In the Tunnel tab, the admin group bit names appear in the Admin Group Include All, Admin Group Include Any, and Admin Group Exclude columns.
4. (Optional) To see the admin group names on the topology map, right-click the blank space in the topology map and select Link Label > TE Admin Group A::Z from the list. The named (colored) links then appear in the topology map.

RELATED DOCUMENTATION

[About the Tunnel Tab | 710](#)

[About the Link Tab | 702](#)

Modify Pathfinder Settings From the Pathfinder CLI

IN THIS SECTION

- [Access the Pathfinder CLI | 199](#)
- [Modify Pathfinder Configuration Settings | 203](#)

The component-specific Pathfinder settings previously maintained in the **northstar.cfg** file are maintained in an internal cache and can be modified by using the Pathfinder CLI. The Pathfinder CLI is very similar to the Junos OS CLI.

NOTE: Certain bootstrap and infrastructure configuration settings continue to be maintained in the **northstar.cfg** file.

If you are not already familiar with the Junos OS CLI, see [Junos OS CLI User Guide](#) which covers:

- Accessing operational and configuration command modes, and switching between modes.
- The concept of command hierarchies.
- Navigation among hierarchy levels.
- Getting help on command syntax including how to display all possible completions for a partial command.
- Helpful keyboard sequences including command completion shortcuts.
- Committing or backing out of configuration changes.

You can also modify the Pathfinder settings from the GUI. See ["Modify Pathfinder Settings From the GUI" on page 206](#).

Access the Pathfinder CLI

To access the Pathfinder CLI:

1. After you successfully install Paragon Pathfinder, log in to the Pathfinder primary server.
2. Launch the Pathfinder CLI:

```
[root@ns]# kubectl exec -it ns-cmgd-<pod-name> -n northstar -c ns-cmgd -- cli
```

<pod-name> is the name of the pod where the Containerized Management Daemon (cMGD) is installed.

To get the cMGD pod name in the cluster, run the following command:

```
[root@ns]# kubectl get pods -n northstar | grep ns-cmgd
```

3. The `>` prompt indicates you are in operational mode. In this mode, you can display the Pathfinder configuration, but you cannot change it. For example:

```
root@ns> show configuration northstar

config-server {
  health-monitor {
    heartbeat-interval 5s;
    poll-interval 10m;
    history-ttl 2d;
    heartbeat-holdown-timer 15s;
  }
}
```

```

}
path-computation-server {
    health-monitor {
        heartbeat-interval 5s;
        poll-interval 10m;
        history-ttl 2d;
        heartbeat-holddown-timer 15s;
    }
}
netconfd {
    in-memory-datastore {
        reconnect-delay 1000;
        reconnect-retries 10000;
    }
}
programmable-rpd-client {
    health-monitor {
        heartbeat-interval 5s;
        poll-interval 10m;
        history-ttl 2d;
        heartbeat-holddown-timer 15s;
    }
    enable-top-prefix-filter;
    publish-top-prefix-only;
}

root@ns1>

```

4. Use the **edit** (or **configure**) command to enter configuration mode (prompt changes to #):

```

root@ns> edit
Entering configuration mode
Users currently editing the configuration:
  root terminal pts/1 (pid 246) on since 2020-08-25 17:26:47 UTC, idle 1d 23:54
    [edit]
  root terminal pts/2 (pid 262) on since 2020-08-26 23:05:31 UTC, idle 17:13:12
    [edit]

[edit]
root@ns1#

```

[edit] indicates the top of the command hierarchy.

5. All Pathfinder configuration commands begin with **set northstar**. Enter **set northstar** with a question mark to display the Pathfinder configuration command top level categories:

```

root@ns1# set northstar ?

Possible completions:
> analytics          General configuration parameters related to analytics
> config-server      Config Server run time parameters
> mladapter          General configuration parameters related to ML Adapter. Common
                    configuration parameters like amqp or database are taken from
                    amqpSettings, but can be overridden for MLAdapter.
> netconfd          General configuration parameters related to netconfd
> path-computation-server Path computation server run time parameters
> peer-engineering   General configuration parameters for EPE and IPE
> programmable-rpd-client General configuration parameters related to the PRPD client
> system
> topology-server    General configuration parameters related to the Topology Server. Common
                    configuration parameters like amqp or database are taken from
                    amqpSettings, but can be overridden for the Topology Server.

[edit]
root@ns1# set northstar

```

6. Continue with any category and a question mark to see the next level breakdown. For example:

```

root@ns1# set northstar config-server ?

Possible completions:
> health-monitor     Configuration parameters related to Health Monitor
+ include-interface-type The interfaces to be published by Config Server
                    Space-separated list enclosed in [ ] or single interface type with
                    no brackets. Indicates
                    discovered interface types to be added to NorthStar.
                    The following interface types are supported:
                    - physical      Physical interfaces : interface name without dot (.)
                    - loopback-mgmt Loopback and management interfaces : interface
                    name starting with lo, fxp, me and em
                    - vrf-if       Interfaces associated with VRF
                    - links-if    Interfaces on links
                    - all         all interfaces

```

```

> log-destination      List of logging configuration
  publish-aslink       Enable ConfigServer to publish aslink created by getipconf to the
Northstar model

[edit]
root@ns1# set northstar config-server

```

7. Continue in this fashion to reach a configuration setting and, if the command requires it, specify the value you wish to change. For example:

```

root@ns1# set northstar config-server include-interface-type ?
Possible completions:
[                Open a set of values
all
links-if
loopback-mgmt
physical
vrf-if
[edit]

root@ns1# set northstar config-server include-interface-type [links-if physical]

[edit]
root@ns1#

```

8. Once you are familiar with the command hierarchy, you can navigate directly to a different level once you are in configuration mode. For example:

```

root@ns> edit
Entering configuration mode
Users currently editing the configuration:
  root terminal pts/0 (pid 162) on since 2020-09-05 17:08:10 UTC, idle 1d 07:59
  [edit]

[edit]
root@ns# edit northstar system health-monitor

[edit northstar system health-monitor]
root@ns# set ?
Possible completions:
  heartbeat-holdown-timer Health monitor holdown timer. Can be expressed as seconds ('s' or
'seconds'). Examples: 30s, 30seconds. (default=15s)

```

```

heartbeat-interval Health monitoring heartbeat interval. Can be expressed as seconds ('s'
or 'seconds'). Examples: 10s, 10seconds. (0 or -ve value = disabled,
default=5s)

history-ttl Health monitor history retention. Can be expressed as days ('d' or
'days'). Examples: 7d, 7days. (default=2d)

poll-interval Health monitor poll interval. Can be expressed as minutes ('m' or
'minutes'). Examples: 4m, 4minutes. (default=10m)

```

```
[edit northstar system health-monitor]
```

Modify Pathfinder Configuration Settings

Table 51 on page 203 lists the Pathfinder configuration settings most likely to require modification. It is not a complete listing of all the available settings.

NOTE: The following configuration settings are not supported currently. We strongly recommend that you do not modify these settings:

- Analytics
- Health Monitor
- Peer Engineering
- Some Topology Server settings: **element-state-mask**, **element-type-mask**, **use-safe-mode**, **do-not-use-davinci-device-model**

Table 51: Frequently-used Pathfinder Configuration Settings

Setting	Command	Description
port (NETCONF)	set northstar netconf-connection-controller device-connection-pool netconf port	Change the default port for NETCONF from 830. In some installations, port 22 is preferred.

Table 51: Frequently-used Pathfinder Configuration Settings (*Continued*)

Setting	Command	Description
exec-pace-rate (EPE)	set northstar peer-engineering egress application exec-pace-rate	Set the maximum rate (in calls per second) at which the egress peer engineering (EPE) Planner executes Pathfinder REST API calls that change the network. NOTE: Currently, this setting is not supported.
target-tag-cookie-range-start (EPE)	set northstar path-computation-server bgp-steering target-tag-cookie-range-start	Set the starting number for path cookie allocation. EPE static routes added by Pathfinder are allocated an unused path cookie starting from the value specified here.
target-tag (EPE)	set northstar path-computation-server bgp-steering target-tag	Set the target prefix for the PRPD filtering community.
polling-interval (Multilayer)	set northstar mladapter polling-interval	Set the polling interval (in seconds) for interfaces without notification support.
lsp-latency-interval (Analytics)	set northstar path-computation-server lsp-latency-interval	Set the interval (in seconds) for which the PCViewer calculates LSP delay and display the data in the Tunnels tab of the GUI (Network > Topology > Tunnels tab > View > Delay).

Table 51: Frequently-used Pathfinder Configuration Settings (*Continued*)

Setting	Command	Description
include-interface-type (all tasks related to interfaces)	set northstar config-server include-interface-type	<p>Set interface types that can be discovered on devices and that are to be used for traffic collection.</p> <p>The supported interface types are:</p> <ul style="list-style-type: none"> • physical: Physical interfaces, expressed as the interface name without a dot (.) in it. • loopback-mgmt: Loopback and management interfaces expressed as the interface name starting with lo, fxp, me, or em. • vrf-if: Interfaces associated with a VRF. • links-if: Interfaces on links. • all: All interfaces. <p>NOTE: configServer publishes to all components only the interface types that you specify. The GUI and data collection only receive information about interfaces representing those interface types.</p> <p>If you modify this setting by clearing interface types that are already represented by interfaces in the Pathfinder model, those existing interfaces remain in the model.</p>

RELATED DOCUMENTATION

| [Paragon Pathfinder Overview](#) | 6

Modify Pathfinder Settings From the GUI

In addition to modifying the settings from the CLI, you can also modify the component-specific Pathfinder settings by using the Paragon Automation GUI.

NOTE: Currently, the Peer Engineering configuration settings are not supported. Although you can modify these settings from the GUI, the changes will not take effect.

To modify the settings from the GUI:

1. Select **Configuration > Network Settings > Pathfinder Settings**.

The Pathfinder page is displayed.

2. Select a component that you want to modify.

The component-specific settings appear on the right side of the page.

3. Modify the settings as needed. [Table 52 on page 206](#) describes the settings that you can modify.

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

4. Click **Save Changes**.

The changes are saved and a confirmation message is displayed.

The modified settings are displayed on the Pathfinder page.

Table 52: Pathfinder Settings

Setting	Description
---------	-------------

Advanced

Table 52: Pathfinder Settings (Continued)

Setting	Description
Sync Network Model	<p data-bbox="548 365 1390 464">Click Sync to refresh the synchronization of the network model. You can use this option if the network model audit has unresolved discrepancies or if the information displayed for the model is out of sync.</p> <p data-bbox="548 495 1341 525">When you sync the network model, this is what happens behind the scenes:</p> <ol data-bbox="548 556 1414 810" style="list-style-type: none"> <li data-bbox="548 556 1414 655">1. Information associated with the network model (nodes, links, LSPs, interfaces, SRLGs, and user-defined parameters) remains intact. Nothing is purged from the database. <li data-bbox="548 686 1328 810">2. The network model is repopulated with live data learned from topology acquisition. <p data-bbox="548 846 1373 911">Table 53 on page 249 describes the effects on various elements in the network when you reset or synchronize the model.</p>

Table 52: Pathfinder Settings (Continued)

Setting	Description
Reset Network Model	<p>WARNING: The Reset Network Model operation is typically appropriate for a lab rather than a production environment. We strongly recommend that you perform this action only under the supervision of Juniper Networks Technical Assistance Center (JTAC).</p> <p>There are two circumstances under which you must reset the network model in order to keep the model in sync with the actual network:</p> <ul style="list-style-type: none"> • The node ISO network entity title (NET) address changes—This can happen when configuration changes are made to support IS-IS. • The routing device's IP address (router ID) changes—The router ID is used by BGP and OSPF to identify the routing device from which a packet originated. The router ID is usually the IP address of the local routing device. If a router ID has not been configured for the device, the IP address of the first interface (of the device) that comes online is used, which is usually the loopback interface. Otherwise, the first hardware interface with an IP address is used. <p>If either of these addresses changes, and you do not perform the Reset Network Model operation, the network model in the Pathfinder database becomes out of sync with the live network.</p> <p>When you reset the network model, this is what happens behind the scenes:</p> <ol style="list-style-type: none"> 1. Information associated with the network model (nodes, links, LSPs, interfaces, SRLGs, and user-defined parameters) is purged from the database (so, we recommend that you do not reset the network model unless it is absolutely required, and that you perform this action only under JTAC supervision). <p>NOTE: Device profiles are not affected.</p> <ol style="list-style-type: none"> 2. The network model is repopulated with live data learned from topology acquisition. <p>Table 53 on page 249 describes the effects on various elements in the network when you reset or synchronize the model.</p>

Table 52: Pathfinder Settings (Continued)

Setting	Description
<i>Settings</i>	<ul style="list-style-type: none"> • <code>publish-aslink</code>—If you enable this toggle button, the ConfigServer publishes the AS link (created by the <code>getipconf</code> command) to Pathfinder. By default, this toggle button is disabled. • <code>update-topology</code>—If you enable this toggle button, the config server adds new nodes and links (not discovered through BGP-LS) in the topology. If you disable (default) this toggle button, new nodes and links aren't added in the topology. This parameter hasn't been tested in all environments and hence, configuring this parameter may not provide the desired results. • <code>include-interface-type</code>—From the list, select one or more interface types that you want the ConfigServer to add to Pathfinder: <ul style="list-style-type: none"> • <code>physical</code> (physical interfaces) • <code>loopback-mgmt</code> (loopback and management interfaces) • <code>vrf-if</code> (interfaces associated with a VRF) • <code>links-if</code> (interfaces on links) • <code>all</code> (all interfaces) <p>NOTE: ConfigServer publishes to all components only the interface types that you specify. The GUI and data collection receive information about the interfaces representing those interface types only. Later, if you modify this setting by clearing the previously selected interface types (that are already represented by interfaces in the Pathfinder model), the information collected for those interfaces remains in the model.</p>

Table 52: Pathfinder Settings (*Continued*)

Setting	Description
Log Destination	<ul style="list-style-type: none"> • name—Displays the name that is used to identify the log destination configuration. • level—From the list, select the severity level of the log messages. The available options are: <ul style="list-style-type: none"> • emergency—The system is unusable. • alert—Immediate action is needed. • critical—Critical condition exists. • error—Error condition. • warning—Warning condition (this is the default value). • notice—Normal but significant condition. • info—Information message. • debug—Debug message. • trace—Trace message. • - (none)—No severity level.
Device Profile Adapter	
<i>Settings</i>	<p>redis-poll-interval—Specify the frequency (in seconds) with which the device profile adapter polls the Redis service for the Redis status.</p> <p>Enter 0 or a negative value to disable this setting.</p> <p>Example: 10s or 10seconds</p> <p>Default: 5s</p>

Table 52: Pathfinder Settings (*Continued*)

Setting	Description
Log Destination	<ul style="list-style-type: none"> • name—Displays the name that is used to identify the log destination configuration. • level—From the list, select the severity level of the log messages. The available options are: <ul style="list-style-type: none"> • emergency—The system is unusable. • alert—Immediate action is needed. • critical—Critical condition exists. • error—Error condition. • warning—Warning condition (this is the default value). • notice—Normal but significant condition. • info—Information message. • debug—Debug message. • trace—Trace message. • - (none)—No severity level.
Mladapter	
<i>Settings</i>	<p>polling-interval—Specify the polling interval (in seconds) for interfaces without notification support. To disable polling, enter 0.</p> <p>Default: 3600</p>

Table 52: Pathfinder Settings (*Continued*)

Setting	Description
Log Destination	<ul style="list-style-type: none"> • name—Displays the name that is used to identify the log destination configuration. • level—From the list, select the severity level of the log messages. The available options are: <ul style="list-style-type: none"> • emergency—The system is unusable. • alert—Immediate action is needed. • critical—Critical condition exists. • error—Error condition. • warning—Warning condition (this is the default value). • notice—Normal but significant condition. • info—Information message. • debug—Debug message. • trace—Trace message. • - (none)—No severity level.

Netconf Connection Controller

Table 52: Pathfinder Settings (Continued)

Setting	Description
Device-Connection Pool	<ul style="list-style-type: none"> • Device-Connection Pool: <ul style="list-style-type: none"> • batch-size—Specify the maximum number of devices (per cycle) for which a NETCONF session must be established with the NETCONF connection controller. Default: 10 • reconnect-delay—Specify the delay (in seconds) after which the NETCONF connection controller attempts to reconnect with the device, when the NETCONF session fails. Default: 30s • disconnect-delay—Specify the delay (in seconds) after which an attempt is made to disconnect the NETCONF session between the NETCONF connection controller and the device, when the session is migrated. Example: 250s or 250seconds Default: 300s • Netconf: <ul style="list-style-type: none"> • keepalive-interval—Specify the interval (in seconds) at which the NETCONF connection controller periodically sends messages to the devices to ensure that the NETCONF session isn't disconnected due to inactivity. Default: 0s • port—Specify a TCP port number to change the default port for NETCONF from 830. In some installations, port 22 is preferred. Default: 830

Table 52: Pathfinder Settings (Continued)

Setting	Description
In-memory-Datastore	<ul style="list-style-type: none"> <li data-bbox="553 373 1390 527"> <p>• connection-pool-size—Specify the maximum number of connections to be maintained, between the microservice and in-memory datastore, in the connection pool.</p> <p>Default: 5</p> <li data-bbox="553 569 1390 722"> <p>• reconnect-delay— Specify the delay (in milliseconds) after which the microservice attempts to reconnect with the in-memory datastore, when the connection between the two fails.</p> <p>Default: 1000</p> <li data-bbox="553 764 1390 917"> <p>• reconnect-retries—Specify the maximum number of times a microservice can attempt to reconnect with the in-memory datastore, when the connection between the two fails.</p> <p>Default: 1000</p>
Infra	<ul style="list-style-type: none"> <li data-bbox="553 1003 1398 1220"> <p>• Cache Memory:</p> <ul style="list-style-type: none"> <li data-bbox="586 1066 1398 1220"> <p>• max-reconnect-attempts—Specify the maximum number of reconnection attempts to be allowed between the NETCONF connection controller and cache-memory, when the connection fails.</p> <p>Default: 30</p> <li data-bbox="553 1262 1414 1661"> <p>• Msg Broker:</p> <ul style="list-style-type: none"> <li data-bbox="586 1325 1414 1478"> <p>• max-queue-length—Specify the maximum allowed length of the queue used by the message broker.</p> <p>Default: 0</p> <p>If you use the default value, the queue is set without any maximum length.</p> <li data-bbox="586 1541 1382 1661"> <p>• prefetch-count—Specify the number of messages to be sent at the same time by the message broker.</p> <p>Default: 200</p>

Table 52: Pathfinder Settings (Continued)

Setting	Description
Network Data Backend	<p>backend-connection—From the list, select either folder or in-memory datastore as the backend datastore from which you want to fetch the topology information:</p> <ul style="list-style-type: none"> • in-memory-datastore—If you select redis, the values that you configured in <i>netconf-connection-controller > in-memory-datastore</i> are applied, and those values are displayed here. • folder-path—If you select file-system, specify the path to the network definition folder. <p>Default: /opt/northstar/data/network_data/</p>
Registry	<ul style="list-style-type: none"> • check-delay—Specify the delay interval (in seconds) between checking registered NETCONF connection controller instances in Redis. <p>Default: 45s</p> <ul style="list-style-type: none"> • publish-delay—Specify the delay interval (in seconds) between attempting registration of NETCONF connection controller instances in Redis. <p>Default: 30s</p>
Rpc Service	<p>workers—Specify the number of threads to be used by the remote procedure call (RPC) service provider.</p> <p>Default: 10</p>

Table 52: Pathfinder Settings (*Continued*)

Setting	Description
Log Destination	<ul style="list-style-type: none"> • name—Displays the name that is used to identify the log destination configuration. • level—From the list, select the severity level of the log messages. The available options are: <ul style="list-style-type: none"> • emergency—The system is unusable. • alert—Immediate action is needed. • critical—Critical condition exists. • error—Error condition. • warning—Warning condition (this is the default value). • notice—Normal but significant condition. • info—Information message. • debug—Debug message. • trace—Trace message. • - (none)—No severity level.
Netconfd	
Device Connection Pool	<p>publish-connected-device-status—Specify the interval (in seconds) for publishing the connection status of all the devices managed by netconfd.</p> <p>Default: 0s</p> <p>If you use the default value, the device connection status is not published.</p>

Table 52: Pathfinder Settings (Continued)

Setting	Description
In-memory Datastore	<ul style="list-style-type: none"> <li data-bbox="548 373 1390 470">• <code>connection-pool-size</code>—Specify the maximum number of connections to be maintained, between the microservice and in-memory datastore, in the connection pool. Default: 5 <li data-bbox="548 569 1390 665">• <code>reconnect-delay</code>—Specify the delay (in milliseconds) after which the microservice attempts to reconnect with the in-memory datastore, when the connection between the two fails. Default: 1000 <li data-bbox="548 764 1390 861">• <code>reconnect-retries</code>—Specify the maximum number of times a microservice can attempt to reconnect with the in-memory datastore, when the connection between the two fails. Default: 1000
Infra	<p data-bbox="548 995 716 1022">Cache Memory:</p> <ul style="list-style-type: none"> <li data-bbox="548 1058 1390 1155">• <code>max-reconnect-attempts</code>—Specify the maximum number of reconnection attempts to be allowed between netconfd and the cache-memory, when the connection fails. Default: 30 <li data-bbox="548 1253 716 1281">• Msg Broker: <ul style="list-style-type: none"> <li data-bbox="586 1316 1419 1377">• <code>max-queue-length</code>—Specify the maximum allowed length of the queue used by the message broker. Default: 0 If you use the default value, the queue is set without any maximum length. <li data-bbox="586 1535 1382 1596">• <code>prefetch-count</code>—Specify the number of messages to be sent at the same time by the message broker. Default: 200

Table 52: Pathfinder Settings (Continued)

Setting	Description
Network Data Backend	<p>backend-connection—Select the backend datastore (either folder or in-memory datastore) from which you want to fetch the topology information:</p> <ul style="list-style-type: none"> • in-memory-datastore—If you select redis, the values that you configured in <i>netconfd > in-memory-datastore</i> are applied, and those values are displayed in this field. • folder-path—If you select file-system, specify the path to the network definition folder. <p>Default: /opt/northstar/data/network_data/</p>
Proxy	<p>workers—Specify the number of threads to be used by the proxy service provider.</p> <p>Default: 10</p>
Registry	<p>check-delay—Specify the delay (in seconds) between checking registered netconfd instances in Redis.</p> <p>Default: 45s</p>

Table 52: Pathfinder Settings (*Continued*)

Setting	Description
Log Destination	<ul style="list-style-type: none"> • name—Displays the name that is used to identify the log destination configuration. • level—From the list, select the severity level of the log messages. The available options are: <ul style="list-style-type: none"> • emergency—The system is unusable. • alert—Immediate action is needed. • critical—Critical condition exists. • error—Error condition. • warning—Warning condition (this is the default value). • notice—Normal but significant condition. • info—Information message. • debug—Debug message. • trace—Trace message. • - (none)—No severity level.
Path Computation Server	

Table 52: Pathfinder Settings (Continued)

Setting	Description
<i>Settings</i>	<ul style="list-style-type: none"> <li data-bbox="553 373 1422 506">• ignore-rro-check—Click the toggle button to enable or disable the Path Computation Server (PCS) to ignore explicit route objects (ERO) and record route objects (RRO) during LSP initial node consolidation. By default, this toggle button is disabled. <li data-bbox="553 541 1422 611">• northstar-vpn—Click the toggle button to enable or disable Pathfinder VPN features in the PCS. By default, this toggle button is disabled. <li data-bbox="553 646 1422 1024">• path-computation-state-timeout—Specify a timeout (in seconds). During stateful path computation for PCEP request (PCReq) and PCEP reply (PCReply), and REST API, the path computation state is held for the specified timeout (rounded to the next multiple of 5). If the user request or the PCC provisions the policy within the timeout, the already computed state and bandwidth are used. If the timeout expires before the user or the PCC provisions the policy, a new path is computed. After the PCC sends the PCReport (PCRpt), the LSP is reported with the computed Explicit Route Object (ERO). For more information, see "Add and Remove LSP Delegation" on page 816. Default: 10 seconds. <li data-bbox="553 1121 1422 1220">• lsp-latency-interval—Specify the frequency (in seconds) at which the PCViewer calculates LSP delay and display the data in the Tunnel tab of the GUI (Network > Topology > Tunnels tab > View > Delay). <li data-bbox="553 1255 1422 1381">• license-check-interval—Specify the interval for the PCS to check for the license (npatpw) file. Default: 3600s <li data-bbox="553 1417 1422 1516">• route-over-logical-SRLink—Click the toggle button to enable or disable the PCS to calculate a route over a logical segment routed (SR) link. By default, this toggle button is disabled. <li data-bbox="553 1551 1422 1650">• SRLG-provisioning—Click the toggle button to enable or disable the provisioning of shared risk link groups (SRLG). By default, this toggle button is disabled. <li data-bbox="553 1686 1422 1785">• provisioning-include-lsp-metric—Click the toggle button to enable or disable the PCS to include the LSP metric when provisioning the LSP. By default, this toggle button is disabled.

Table 52: Pathfinder Settings (Continued)

Setting	Description
	<p>NOTE: The PCEP reports the route metric for Juniper routers. For more information, see Basic LSP Configuration. You can set the route metric for an LSP and configure the global preference of Paragon Pathfinder from the Configuration > Network Settings > Pathfinder Settings > Path Computation Server page. However, for non-Juniper routers, the reported metric is the sum of the IGP metrics on all outgoing interfaces along a particular path from the source to the destination.</p> <ul style="list-style-type: none"> • round-trip-delay-on-remote-end—Click the toggle button to enable or disable (default) configuring the link latency at the remote-end router as the same value as the link latency at the local-end router. Link latency is equal to half the value of the measured delay, which is calculated as the RTT in one direction. • disable-ecmp-tree-calculation—Click the toggle button to enable or disable the calculation of equal-cost multi-paths (ECMP). If you enable this toggle button, equal-cost multi-paths (ECMP) are not calculated. If you disable (default) this toggle button, ECMPs are calculated; if paths to a remote destination have the same cost, then traffic is distributed between them in equal proportion. • disable-node-sid-calculation—Click the toggle button to enable or disable node segment ID (SID) calculation. If you enable this toggle button, SIDs for the nodes are not calculated. If you disable (default) this toggle button, segment IDs (SID) for the nodes are calculated. • lsp-provision-queue-size—Specify the size for the LSP provisioning queue. Default: 50 • pcep-speaker-id—Specify the PCEP speaker ID to be used for P2MP flow mapping. Default: northstar • zero-bandwidth-signalling—Click the toggle button to enable or disable the zero-bandwidth signaling feature. By default, this toggle button is disabled. • ecmp-placement-method—Specify one of the following options as the ECMP placement method. After considering link metrics, if there are multiple ECMPs available, Pathfinder selects a path based on the value you specify here. If you don't specify a value, the first random path is selected for all the LSPs. Random—Pathfinder randomly selects one of the ECMP paths.

Table 52: Pathfinder Settings (Continued)

Setting	Description
	<p>Least Fill—Pathfinder selects the path with the maximum available bandwidth.</p> <p>Most Fill—Pathfinder selects the path with the minimum available bandwidth.</p> <ul style="list-style-type: none"> monitor-mode—Click the toggle button to enable or disable the monitor mode in the PCS. By default, this toggle button is disabled. <p>If you enable monitor mode, the PCS doesn't send LSP provisioning orders to the PCC. You use this option when you want to monitor the topology but don't want the LSPs provisioned. If you disable monitor mode, the PCS sends LSP provisioning orders to the PCC.</p>
Analytics	<ul style="list-style-type: none"> link-utilization-threshold—Specify the threshold value for link utilization. When traffic on a link exceeds this value, the controller triggers rerouting for label-switched paths (LSPs). If a threshold isn't specified, the LSP is not rerouted. If 0 is specified, links are blocked. Range: 0 through 100. packet-loss-threshold—Specify the threshold value for packet loss on all links. If the packet loss on a link exceeds this value, the link is considered unstable, and the PCS triggers a maintenance event on the link. If you don't specify a value or specify 0, the PCS doesn't act in case of packet loss. Range: 0 through 100. reroute-minimum-interval—Specify the minimum interval (in minutes or m) after which the controller reacts to any traffic or delay violations. If you don't specify an interval, the LSP is not rerouted in case of a violation. Range: 1m through 300m.

Table 52: Pathfinder Settings (Continued)

Setting	Description
BGP Steering	<ul style="list-style-type: none"> <li data-bbox="553 369 1370 432">• <code>target-tag</code>—Specify the target prefix for the programmable routing protocol process (PRPD) filtering community. Default: 42 <li data-bbox="553 531 1370 594">• <code>target-tag-cookie-range-start</code>—Set the starting number for allocating path cookies. EPE static routes added by Pathfinder are allocated an unused path cookie starting from the value specified here. Default: 42 <li data-bbox="553 730 1403 898">• <code>steering-route-preference</code>—Specify the relative preference value for the PRPD steering route. If you specify a sign (+ or -) for the value, the preference is considered the relative increase or decrease of the preference for static routes (relative to BGP routes). If you don't specify a sign, the value is considered an absolute value. Range: -127 through 127 Default: 1 <li data-bbox="553 1056 1386 1161">• <code>steering-aggregate-colors</code>—Click the toggle button to enable the creation of a steering entry (per ingress prefix) that contains the route target with all the colors. By default, this toggle button is disabled. NOTE: This parameter hasn't been tested in all environments and hence, configuring this parameter may not provide the desired results.
In Memory Datastore	<ul style="list-style-type: none"> <li data-bbox="553 1325 1370 1461">• <code>disable-pipeline</code>—If you enable this toggle button, the in-memory datastore pipeline feature is disabled (meaning, multiple entries are written to the in-memory datastore one by one). If you disable this toggle button, multiple entries are written to the in-memory datastore in a single operation. <li data-bbox="553 1503 1354 1566">• <code>pipeline-threshold</code>—Specify the buffer threshold value for the in-memory datastore pipeline. Default: 65536

Table 52: Pathfinder Settings (Continued)

Setting	Description
Link Flap	<ul style="list-style-type: none"> • flap-interval—Specify the interval (in seconds) after which the link flap count is reset. If a link remains in the same status (Up or Down) longer than this interval, the counter is reset and the link is no longer considered flapped. Range: 1 through 300s • flap-count—Specify the maximum value for the link flap count. When a link goes from Up to Down, the Path Computation Element (PCE) increments the counter on that link. When the counter reaches the maximum link flap count, the link is considered flapped. Flapped links carry a large penalty, so are not preferred by the PCS.
LSP To Path Computation Instance	<ul style="list-style-type: none"> • LSP Request Discriminator SR NodeSID: Instance-type—From the list, select the instance of the PCS that manages the segment routing (SR) LSPs tagged with Use node SIDs: <ul style="list-style-type: none"> • default—To choose the default PCS. • SRPCServer—To choose the SR PCS with anycast ID support. • LSP Request Discriminator SR Test: From the list, select the instance of the PCS that manages the SR LSPs: <ul style="list-style-type: none"> • default—To choose the default PCS. • SRPCServer—To choose the SR PCS.
Segment Routing Policy	<p>segment-routing-policy:</p> <ul style="list-style-type: none"> • originator-asn—Specify the originator Autonomous System Number (ASN) from which the segment routing policy originates. Default: 0 • originator-ip—Specify the IP address from where the segment routing policy originates.

Table 52: Pathfinder Settings (Continued)

Setting	Description
Log Destination	<ul style="list-style-type: none"> • name—Displays the name that is used to identify the log destination configuration. • level—From the list, select the severity level of the log messages. The available options are: <ul style="list-style-type: none"> • emergency—The system is unusable. • alert—Immediate action is needed. • critical—Critical condition exists. • error—Error condition. • warning—Warning condition (this is the default value). • notice—Normal but significant condition. • info—Information message. • debug—Debug message. • trace—Trace message. • - (none)—No severity level.
Path Optimization	
Timer	<p>Specify the frequency (in minutes) with which path optimization is triggered automatically.</p> <p>NOTE: The optimization is based on the current network, and not on the most recent Path Analysis report.</p>
Programmable RPD Client	

Table 52: Pathfinder Settings (Continued)

Setting	Description
<i>Settings</i>	<ul style="list-style-type: none"> • retry-interval—Specify the frequency (in seconds) with which the programmable routing protocol process (PRPD) client retries to connect with the PRPD daemon (on the router) after the connection between the two fails. Example: 60s or 60seconds Default: 30s • top-prefix-filter—Specify how the top prefix filter is applied: <ul style="list-style-type: none"> • Onwards—Only the routes learned in the future will be checked against the filter and published if the routes meet the filter criteria; existing routes are retained. • Immediately—Existing routes are cleared and all routes (existing and those learned in the future) are checked against the filter and published only if the routes meet the filter criteria.
Log Destination	<ul style="list-style-type: none"> • name—Displays the name that is used to identify the log destination configuration. • level—From the list, select the severity level of the log messages. The available options are: <ul style="list-style-type: none"> • emergency—The system is unusable. • alert—Immediate action is needed. • critical—Critical condition exists. • error—Error condition. • warning—Warning condition (this is the default value). • notice—Normal but significant condition. • info—Information message. • debug—Debug message. • trace—Trace message. • - (none)—No severity level.

Table 52: Pathfinder Settings (Continued)

Setting	Description
System	
In Memory Datastore	<ul style="list-style-type: none"> <li data-bbox="553 470 1386 569">• <code>connection-pool-size</code>—Specify the maximum number of connections to be maintained, between the microservice and in-memory datastore, in the connection pool. Default: 5 <li data-bbox="553 667 1386 766">• <code>reconnect-delay</code>—Specify the delay (in milliseconds) after which the microservice attempts to reconnect with the in-memory datastore, when the connection between the two fails. Default: 1000 <li data-bbox="553 865 1386 963">• <code>reconnect-retries</code>—Specify the maximum number of times a microservice can attempt to reconnect with the in-memory datastore, when the connection between the two fails. Default: 1000
Messaging Bus	<ul style="list-style-type: none"> <li data-bbox="553 1100 1386 1199">• <code>reconnect-retries</code>—Specify the maximum number of times a microservice can attempt to reconnect with the messaging bus, when the connection between the two fails; 0 indicates no attempts and -1 indicates infinite retries. Default: -1 <li data-bbox="553 1297 1386 1396">• <code>reconnect-delay</code>—Specify the delay (in milliseconds) after which the microservice attempts to reconnect with the messaging bus, when the connection between the two fails. Default: 1000 <li data-bbox="553 1495 1386 1593">• <code>max-channels</code>—Specify the maximum number of channels that can be multiplexed over a single connection between the messaging bus and the microservice. Default: 128

Table 52: Pathfinder Settings (Continued)

Setting	Description
Persistent Datastore	<ul style="list-style-type: none"> • Persistent datastore settings: <ul style="list-style-type: none"> • connection-pool-size—Specify the maximum number of connections to be maintained, between the persistent datastore and the microservice, in the connection pool. Default: 5 • reconnect-delay— Specify the delay (in milliseconds) after which the microservice attempts to reconnect with the persistent datastore, when the connection between the two fails. Default: 1000 • Log Destination: <ul style="list-style-type: none"> • name—Specify a unique name for the log destination configuration. • level—From the list, select the severity level of the log messages. The available options are: <ul style="list-style-type: none"> • emergency—The system is unusable. • alert—Immediate action is needed. • critical—Critical condition exists. • error—Error condition. • warning—Warning condition (this is the default value). • notice—Normal but significant condition. • info—Information message. • debug—Debug message. • trace—Trace message. • - (none)—No severity level.

Table 52: Pathfinder Settings (Continued)

Setting	Description
Scheduler	<p>device-profile-update-interval—Specify the interval (in milliseconds) after which the internal timer in the scheduler polls for device profile updates.</p> <p>Default: 5000</p>
Scheduler > Distributed Task Queue	<ul style="list-style-type: none"> thread-pool-size—Specify the number of threads that can be run for the Device Collection task. <p>Range: 1 through 255</p> <p>Default: 10</p> collector-max-pool-size—Specify the maximum size (in bytes) of queues to be allowed in the distributed task collector messaging bus. <p>Default: 524288000</p>
Scheduler > Tasks	
Collection Cleanup	<ul style="list-style-type: none"> rollup-data-retention-duration—Specify the retention period (in days) for aggregated data. Enter 0d to disable the retention of aggregated data. <p>Example: 800d or 800days</p> <p>Default: 180d</p> raw-data-retention-duration—Specify the retention period (in days) for raw data logs. Enter 0d to disable the retention of raw data logs. <p>Example: 10days or 10d</p> <p>Default: 14d</p> interval—Specify how often (in days) the collection cleanup task is executed. Enter 0 to disable the cleanup task. <p>Example: 7d or 7days</p> <p>Default: 1d</p>

Table 52: Pathfinder Settings (Continued)

Setting	Description
Demand Reports	<p>demand-reports:</p> <p>as-demand-bucket-size—Specify the maximum number of autonomous system (AS) demand report entries that can be stored in the bucket. This parameter isn't required currently but may be required in future releases.</p> <p>Default: 100</p>
Device Collection	<ul style="list-style-type: none"> • data-path—Specify the path to the directory for storing the data collected by the Device Collection task. Default: /opt/northstar/data • enable-ptalk-logging—If you enable this toggle button, a ptalk log file is generated for each device that is part of the device collection task, when the task is run. You can view the log files from the /opt/pcs/log/ directory in the ptalkserver container and use the logs for debugging. If you disable (default) this toggle button, the ptalk log file isn't generated. • timeout—Specify the time (in seconds) after which the device collection task times out. Example: 3000s or 3000seconds Default: 1800s • enable-update-device-profile—If you enable this toggle button, the device-type attribute is updated in device profiles as part of the Device Collection task. If you disable (default) this toggle button, the device-type attribute is not updated. • ptalk-timeout—Specify the time (in seconds) until which the ptalk server waits for a response (to the ptalk request) from each device. If the server doesn't receive a response from a device within the specified time, the ptalk request for that device times out. You must then increase the timeout period and re-run the device collection task for those devices. Default: 600s. • purge-cutoff—Specify the purge cut-off period (in days). The device collection task deletes all files that are older than the cut-off period specified here. Example: 10d or 10days Default: 7d

Table 52: Pathfinder Settings (Continued)

Setting	Description
File Transfer	<ul style="list-style-type: none"> interval—Specify the interval (in seconds) at which the system transfers files from the active node to other cluster nodes. Example: 3000s or 3000seconds Default: 3600s paths—Specify the absolute directory path to one or more directories in other cluster nodes for transferring files from the active node.
Rollup	<ul style="list-style-type: none"> interval—Specify how often the ESRollup system task is run (in hours). The ESRollup system task executes the esrollup.py script to aggregate the previous hour's data. The ESRollup task is called from the Pathfinder server. You can view (but not modify) the rollup task on the Task Scheduler page (Administration > Task Scheduler). Example: 3h or 3hours Default: 1h NOTE: We recommend that you do not change this default value except to disable aggregation. If you want to disable data aggregation, set the value to 0h. max-worker-process—Specify the maximum number of worker processes that can be used for the ESRollup system task. Default: 4 or equivalent to the number of CPUs (whichever value is smaller) bulk-insert-record-count—Specify the maximum number of records in bulk, which can be inserted to the PostgreSQL database in a single operation. Default: 5000

Table 52: Pathfinder Settings (Continued)

Setting	Description
Scheduled LSP Sizing	<ul style="list-style-type: none"> • stats-collection-interval—Specify the interval (in seconds or minutes) at which the native sensor collects statistics related to the LSP. Example: 30s or 1m Default: 1m • Bandwidth Sizing: <ul style="list-style-type: none"> • stats-query-workers—Specify the number of threads that must be used to query the statistics from the native datastore. Default: 1 • Log Destination: <ul style="list-style-type: none"> • name—Displays the name that is used to identify the log destination configuration. • level—From the list, select the severity level of the log messages. The available options are: <ul style="list-style-type: none"> • emergency—The system is unusable. • alert—Immediate action is needed. • critical—Critical condition exists. • error—Error condition. • warning—Warning condition (this is the default value). • notice—Normal but significant condition. • info—Information message. • debug—Debug message. • trace—Trace message.

Table 52: Pathfinder Settings (Continued)

Setting	Description
	<ul style="list-style-type: none"> • - (none)—No severity level. • Container LSP Normalization > Log Destination: <ul style="list-style-type: none"> • name—Displays the name that is used to identify the log destination configuration. • level—From the list, select the severity level of the log messages. The available options are: <ul style="list-style-type: none"> • emergency—The system is unusable. • alert—Immediate action is needed. • critical—Critical condition exists. • error—Error condition. • warning—Warning condition (this is the default value). • notice—Normal but significant condition. • info—Information message. • debug—Debug message. • trace—Trace message. • - (none)—No severity level.
SNMP Collection	<ul style="list-style-type: none"> • timeout—Specify the time (in seconds) after which the SNMP polling task times out. Example: 10s or 10seconds Default: 3s • retries—Specify the maximum number of retries to be allowed for an SNMP poll. Default: 3

Table 52: Pathfinder Settings (Continued)

Setting	Description
Log Destination	<ul style="list-style-type: none"> • name—Specify a unique name for the log destination configuration. • level—From the list, select the severity level of the log messages. The available options are: <ul style="list-style-type: none"> • emergency—The system is unusable. • alert—Immediate action is needed. • critical—Critical condition exists. • error—Error condition. • warning—Warning condition (this is the default value). • notice—Normal but significant condition. • info—Information message. • debug—Debug message. • trace—Trace message. • - (none)—No severity level.
Topology Filter	
Data Persistence	<ul style="list-style-type: none"> • persistent-storage-init-wait—Specify the wait time (in seconds) for a connection to be established between the topology filter and persistent storage before exiting with an initialization failure. Default: 45s • file-store-path—Specify the default file path for storing output data of the topology filter if persistent storage isn't available. Default: /opt/northstar/data

Table 52: Pathfinder Settings (*Continued*)

Setting	Description
Messaging Bus	<ul style="list-style-type: none"> <li data-bbox="553 373 1393 470">• <code>reconnect-retries</code>—Specify the maximum number of times a microservice can attempt to reconnect with the messaging bus, when the connection between the two fails. 0 indicates no attempts and -1 (default) indicates infinite retries. <li data-bbox="553 510 1317 606">• <code>reconnect-delay</code>—Specify the delay (in milliseconds) after which the microservice attempts to reconnect with the messaging bus, when the connection between the two fails. Default: 1000 <li data-bbox="553 705 1349 802">• <code>max-channels</code>—Specify the maximum number of channels that can be multiplexed over a single connection between the messaging bus and the microservice. Default: 128

Table 52: Pathfinder Settings (Continued)

Setting	Description
Persistent Datastore	<ul style="list-style-type: none"> • Persistent datastore settings: <ul style="list-style-type: none"> • connection-pool-size—Specify the maximum number of connections to be maintained, between persistent datastore and the microservice, in the connection pool. Default: 5 • reconnect-delay—Specify the delay (in milliseconds) after which the microservice attempts to reconnect with the persistent datastore, when the connection between the two fails. Default: 5000 • Log Destination: <ul style="list-style-type: none"> • name—Specify a unique name for the log destination configuration. • level—From the list, select the severity level of the log messages. The available options are: <ul style="list-style-type: none"> • emergency—The system is unusable. • alert—Immediate action is needed. • critical—Critical condition exists. • error—Error condition. • warning—Warning condition (this is the default value). • notice—Normal but significant condition. • info—Information message. • debug—Debug message. • trace—Trace message. • - (none)—No severity level.

Table 52: Pathfinder Settings (*Continued*)

Setting	Description
Log Destination	<ul style="list-style-type: none"> • name—Displays the name that is used to identify the log destination configuration. • level—From the list, select the severity level of the log messages. The available options are: <ul style="list-style-type: none"> • emergency—The system is unusable. • alert—Immediate action is needed. • critical—Critical condition exists. • error—Error condition. • warning—Warning condition (this is the default value). • notice—Normal but significant condition. • info—Information message. • debug—Debug message. • trace—Trace message. • - (none)—No severity level.
Topology Server	

Table 52: Pathfinder Settings (Continued)

Setting	Description
Application	<ul style="list-style-type: none"> <li data-bbox="553 373 1406 470"> <p>• use-live-rsvp-bw-over-configured—Click the toggle button to enable the usage of live RSVP bandwidth (obtained from the live network) instead of the RSVP bandwidth that you configured.</p> <p>By default, this toggle button is disabled.</p> <li data-bbox="553 569 1406 665"> <p>• use-nokia-path-workaround—Click the toggle button to append a double colon (::) between the device name and routing path name strings for Nokia devices.</p> <p>By default, this toggle button is disabled.</p> <li data-bbox="553 730 1406 884"> <p>• source-of-truth—Click the toggle button to modify network topology (add, modify, or delete nodes, links, and tunnels) when one of the deployment is down.</p> <p>Default: Disabled</p> <li data-bbox="553 928 1406 1081"> <p>• sync-topology-after-failure—Click the toggle button in the non-source-of-truth cluster to synchronize the topology without having to restart the toposerver pod when recovering from missing beacon in a disaster recovery deployment.</p> <p>Default: Disabled</p> <li data-bbox="553 1125 1406 1278"> <p>• toposerver-beacon-interval—Enter the frequency (in seconds) with which the topology servers in different Paragon Automation deployment exchange beacons in a disaster recovery deployment.</p> <p>Default: 5s.</p> <li data-bbox="553 1323 1406 1535"> <p>• use-unnumbered-interface-workaround—Click the toggle button to enable the usage of interface index (IfIndex) to correlate link events for unnumbered interfaces.</p> <p>If you enable this toggle button, IfIndex is not used for correlating link events.</p> <p>By default, this toggle button is disabled.</p> <li data-bbox="553 1579 1406 1696"> <p>• use-prefix-link-matching—Click the toggle button to enable the usage of prefixes to correlate bidirectional links events.</p> <p>By default, this toggle button is disabled.</p>

Table 52: Pathfinder Settings (Continued)

Setting	Description
	<ul style="list-style-type: none"><li data-bbox="548 359 1401 491">• do-not-publish-rest-updates—If you enable this toggle button, the topology server doesn't publish object updates to the REST server. If you disable (default) this toggle button, the topology server will publish object updates to the REST server.<li data-bbox="548 531 1417 632">• do-not-suppress-beacon-message—If you enable this toggle button, the beacon messages will be written to log files. If you disable (default) this toggle button, the beacon messages will be suppressed.<li data-bbox="548 669 1417 835">• no-netconf-pathname-allocation—If you enable this toggle button, the topology server doesn't check for path names allocated to LSPs configured through NETCONF and also doesn't set the path name for the planned LSP. If you disable (default) this toggle button, the topology server will check the path names and set the path name for the planned LSP.

Table 52: Pathfinder Settings (Continued)

Setting	Description
Application Tuning	<ul style="list-style-type: none"> <li data-bbox="553 373 1409 506"> <p>• <code>pce-beacon-holddown</code>—Specify the time (in seconds) within which the topology server must receive a beacon from the PCE server. If the PCE beacon is not received within the specified time, the connection between the topology server and the PCEP server is marked as down.</p> <p>Example: 80s or 80 seconds</p> <p>Default: 90s</p> <li data-bbox="553 663 1409 894"> <p>• <code>resync-unresolved-node-threshold</code>—Specify a threshold value for unresolved objects.</p> <p>During the discovery and updation of the topology, the topology server maintains a list of network objects that cannot be resolved to nodes. If the number of unresolved objects exceeds the threshold value specified here, a resynchronization is automatically triggered with the BGP-LS topology source.</p> <p>Default: 10</p> <li data-bbox="553 993 1409 1056"> <p>• <code>lsp-topo-sync-timeout</code>—Specify the frequency (in seconds) with which the topology server synchronizes the PCEP LSP database globally.</p> <p>Default: 120s</p> <li data-bbox="553 1155 1409 1308"> <p>• <code>pce-restart-holddown</code>—Specify the maximum time (in seconds) until which the topology server waits, after the PCEServer starts, to request the PCEP LSP database.</p> <p>Default: 30s</p> <li data-bbox="553 1352 1409 1505"> <p>• <code>message-queue-low-watermark</code>—Specify the low watermark for the messaging queue. If the number of elements in the messaging queue reaches the value specified here, the topology server resumes LSP provisioning.</p> <p>Default: 5000</p> <li data-bbox="553 1549 1409 1738"> <p>• <code>message-queue-high-watermark</code>—Specify the high watermark for the messaging queue. If the number of elements in the messaging queue reaches the value specified here, the topology server pauses LSP provisioning until the number of elements in the messaging queue reaches the low watermark.</p> <p>Default: 500000</p>

Table 52: Pathfinder Settings (Continued)

Setting	Description
	<ul style="list-style-type: none"> message-queue-overflow-watermark—Specify the overflow watermark for the messaging queue. If the number of elements in the messaging queue reaches the value specified here, the topology server purges the messaging queue. Default: 0 (meaning no overflow watermark is set and thus, purging is disabled.)
	<ul style="list-style-type: none"> pce-watermark-interval—Specify the frequency (in seconds) with which the topology server checks the high and low watermark rates for PCEP messages. Default: 10s
	<ul style="list-style-type: none"> pce-high-watermark-timeout—Specify the maximum time (in milliseconds) that the topology server waits for the PCEP message rate to reach the low watermark, after which, it resumes the processing of PCEP messages. Default: 60000
	<ul style="list-style-type: none"> pce-high-watermark-rate—Specify the PCEP message rate to trigger the high watermark. If the PCEP message rate reaches the value specified here, the topology server pauses the processing of PCEP messages. Default: -1.0 (to disable the triggering of low watermark)
	<ul style="list-style-type: none"> pce-high-watermark-count-threshold—Specify the number of times the PCEP message rate must reach the high watermark, after which, it exceeds the high watermark threshold. If the count exceeds the threshold specified here, the counters for the received messages and for watermark are reset, and the pce-high-watermark-timeout is updated. Default: 3
	<ul style="list-style-type: none"> pce-low-watermark-rate—Specify the PCEP message rate to trigger the low watermark. If low watermark is triggered, the topology server resumes the processing of PCEP messages. Default: -1.0 (to disable the triggering of low watermark)
	<ul style="list-style-type: none"> pce-low-watermark-count-threshold—Specify the number of times the PCEP message rate must reach the low watermark, after which, it exceeds the low watermark threshold. If the count exceeds the threshold specified here, the counters for the received messages and for watermark are reset. Default: 3

Table 52: Pathfinder Settings (Continued)

Setting	Description
Data Persistence	<ul style="list-style-type: none"> <li data-bbox="553 373 1409 506"> <p>• <code>persistent-storage-init-wait</code>—Specify the maximum time (in seconds) until which the topology server can wait for a connection to be established with the persistent datastore. If the wait time exceeds the specified value, the topology server exits the topology server application with an initialization failure.</p> <p>Example: 10s or 10seconds</p> <p>Default: 45s</p> <li data-bbox="553 663 1409 785"> <p>• <code>network-snapshot-store-path</code>—Specify the output directory in which you want to save a snapshot of the network topology.</p> <p>Default: <code>/opt/northstar/data/network_archive/NorthStar</code></p> <li data-bbox="553 825 1338 911"> <p>• <code>file-store-path</code>—Specify the default directory for storing event data files.</p> <p>Default: <code>/opt/northstar/data</code></p> <li data-bbox="553 951 1354 1073"> <p>• <code>debug-file-store-path</code>—Specify the default directory for storing the debug output files.</p> <p>Default: <code>/opt/northstar/data</code></p> <li data-bbox="553 1113 1414 1289"> <p>• <code>network-snapshot-store-interval</code>—Specify the maximum time (in seconds) until which the network topology snapshot can be stored in the persistent datastore.</p> <p>Example: 3000s or 3000seconds</p> <p>Default: 3600s</p> <li data-bbox="553 1329 1382 1451"> <p>• <code>persist-lsp-topology-object</code>—Click the toggle button to enable or disable the persistence of the LSP topology object in the persistent datastore.</p> <p>By default, this toggle button is disabled.</p> <li data-bbox="553 1491 1377 1587"> <p>• <code>do-not-persist-pcep-lsp-events</code>—If you enable this toggle button, PCEP LSP events are not stored in the persistent datastore. If you disable (default) this toggle button, PCEP LSP events are stored in the persistent datastore.</p> <li data-bbox="553 1627 1409 1749"> <p>• <code>persist-pcep-pcc-node-events</code>—Click the toggle button to enable or disable the persistence of PCEP node events in the persistent datastore.</p> <p>By default, this toggle button is disabled.</p>

Table 52: Pathfinder Settings (Continued)

Setting	Description
	<ul style="list-style-type: none"> <li data-bbox="553 359 1414 491">• do-not-persist-topology-advertisements—If you enable this toggle button, BGP-LS advertisements are not stored in the persistent datastore. If you disable (default) this toggle button, BGP-LS advertisements are stored in the persistent datastore. <li data-bbox="553 533 1414 665">• do-not-persist-provision-requests—If you enable this toggle button, provisioning requests are not stored in the persistent datastore. If you disable (default) this toggle button, provisioning requests are stored (in JSON format) in the persistent datastore. <li data-bbox="553 707 1414 840">• do-not-persist-lsp-events—If you enable this toggle button, LSP events (event description and object data) events are not stored in the persistent datastore. If you disable (default) this toggle button, LSP events are stored in the persistent datastore. <li data-bbox="553 882 1414 1014">• do-not-persist-node-events—If you enable this toggle button, node events (event description and object data) are not stored in the persistent datastore. If you disable (default) this toggle button, node events are stored in the persistent datastore. <li data-bbox="553 1056 1414 1188">• do-not-persist-link-events—If you enable this toggle button, link events (event description and object data) are not stored in the persistent datastore. If you disable (default) this toggle button, link events are stored in the persistent datastore. <li data-bbox="553 1230 1414 1365">• do-not-persist-lsps—If you enable this toggle button, LSP data is not stored in the persistent datastore. If you disable (recommended) this toggle button, LSP data is stored in the persistent datastore. NOTE: We recommend that you do not disable the persistence of LSP data (that is, we recommend that you do not enable the toggle button). <li data-bbox="553 1512 1414 1644">• persist-topology-snapshot—Click the toggle button to enable or disable the persistence of the topology snapshot in the persistent datastore. By default, this toggle button is disabled. <li data-bbox="553 1686 1414 1820">• persist-demand-events—Toggle the button to enable or disable (default) the persistence of demand events in the persistent datastore. NOTE: For demand events to be persistent, LSP events should be persistent.

Table 52: Pathfinder Settings (*Continued*)

Setting	Description
	<ul style="list-style-type: none"> persist-topology-objects-to-file—Click the toggle button to enable or disable saving the BGP Monitoring Protocol (BMP) events to a file in binary format. By default, this toggle button is disabled.
In-memory Datastore	<p>in-memory-datastore:</p> <ul style="list-style-type: none"> connection-pool-size—Specify the maximum number of connections to be maintained, between the microservice and in-memory datastore, in the connection pool. Default: 5 reconnect-delay—Specify the delay (in milliseconds) after which the microservice attempts to reconnect with the in-memory datastore, when the connection between the two fails. Default: 1000 reconnect-retries—Specify the maximum number of times a microservice can attempt to reconnect with the in-memory datastore, when the connection between the two fails. Default: 1000

Table 52: Pathfinder Settings (Continued)

Setting	Description
Messaging Bus	<ul style="list-style-type: none"> <li data-bbox="553 373 1393 470">• <code>reconnect-retries</code>—Specify the maximum number of times a microservice can attempt to reconnect with the messaging bus, when the connection between the two fails. 0 indicates no attempts and -1 (default) indicates infinite retries. <li data-bbox="553 510 1317 606">• <code>reconnect-delay</code>—Specify the delay (in milliseconds) after which the microservice attempts to reconnect with the messaging bus, when the connection between the two fails. Default: 1000 <li data-bbox="553 705 1349 802">• <code>max-channels</code>—Specify the maximum number of channels that can be multiplexed over a single connection between the messaging bus and the microservice. Default: 128 <li data-bbox="553 900 1398 997">• <code>use-federated-exchange</code>—Click this toggle button to synchronize the topology between two clusters in a disaster-recovery deployment. After changing this setting, you must restart the toposerver pod for the change to take effect. Default: Disabled.

Table 52: Pathfinder Settings (Continued)

Setting	Description
Persistent Datastore	<ul style="list-style-type: none"> • Persistent datastore settings: <ul style="list-style-type: none"> • connection-pool-size—Specify the maximum number of connections to be maintained, between persistent datastore and the microservice, in the connection pool. Default: 5 • reconnect-delay—Specify the delay (in milliseconds) after which the microservice attempts to reconnect with the persistent datastore, when the connection between the two fails. Default: 5000 • Log Destination: <ul style="list-style-type: none"> • name—Specify a unique name for the log destination configuration. • level—From the list, select the severity level of the log messages. The available options are: <ul style="list-style-type: none"> • emergency—The system is unusable. • alert—Immediate action is needed. • critical—Critical condition exists. • error—Error condition. • warning—Warning condition (this is the default value). • notice—Normal but significant condition. • info—Information message. • debug—Debug message. • trace—Trace message. • - (none)—No severity level.

Table 52: Pathfinder Settings (Continued)

Setting	Description
Topology Acquisition	<ul style="list-style-type: none"> <li data-bbox="553 373 1398 541"> <p>• <code>retry-before-exit-count</code>—Specify the number of times the topology server can request a topology refresh (after it connects to a BMP topology source) to obtain the topology that contains the links. If the topology server doesn't receive the topology even after the specified refresh requests, it exits the topology server application.</p> <p>Default: 0 (meaning infinite retries)</p> <li data-bbox="553 642 1398 737"> <p>• <code>retry-delay</code>—Specify the maximum time (in seconds) that the topology server can wait before requesting a topology refresh again (after it connects to a BMP topology source).</p> <p>Example: 30s or 30seconds</p> <p>Default: 5s</p> <li data-bbox="553 898 1398 993"> <p>• <code>reconnection-count</code>—Specify the number of times the topology server can try to reconnect to the BGP-LS topology source before exiting the topology server application.</p> <p>Default: 0 (meaning infinite retries)</p> <li data-bbox="553 1094 1398 1188"> <p>• <code>reconnection-delay</code>—Specify the maximum time (in seconds) that the topology server can wait before attempting to reconnect to the BGP-LS topology source.</p> <p>Example: 30s or 30seconds</p> <p>Default: 5s</p> <li data-bbox="553 1314 1398 1409"> <p>• <code>refresh-holddown</code>—Specify the maximum time (in seconds) that the topology server can wait before requesting for a topology refresh after it connects to a BMP topology source but doesn't receive the topology.</p> <p>Example: 100s or 100seconds</p> <p>Default: 300s</p> <li data-bbox="553 1570 1398 1738"> <p>• <code>eor-timeout</code>—: Specify the maximum time (in seconds) that the topology server can wait to receive the topology from the BGP-LS topology source. After the complete topology is sent, the topology source sends an end-of-Routing Information Base (end-of-RIB or EOR) message that indicates the completion of topology update. If the topology server doesn't receive this EOR message</p>

Table 52: Pathfinder Settings (Continued)

Setting	Description
	<p>within the time that you specify here, an EOR timeout is triggered and the topology server sends another topology refresh request.</p> <p>Example: 15s or 15seconds</p> <p>Default: 20s</p>
Log Destination	<ul style="list-style-type: none"> • name—Displays the name that is used to identify the log destination configuration. • level—From the list, select the severity level of the log messages. The available options are: <ul style="list-style-type: none"> • emergency—The system is unusable. • alert—Immediate action is needed. • critical—Critical condition exists. • error—Error condition. • warning—Warning condition (this is the default value). • notice—Normal but significant condition. • info—Information message. • debug—Debug message. • trace—Trace message. • - (none)—No severity level.

[Table 53 on page 249](#) lists the effects of resetting or synchronizing the network model.

Table 53: Effects of Resetting or Synchronizing the Network Model

Element	Is the element removed from the database?		Is the item sent back to the controller by the live network?		Could data be lost?	
	Reset	Sync	Reset	Sync	Reset	Sync
IP nodes	Yes	No	Yes	Yes	Yes for design attributes, such as user-defined node name	No
IP links	Yes	No	Yes	Yes	Yes for design attributes such as Comment	No
PCC-controlled LSPs	Yes	No	Yes	Yes	No	No
PCC-delegated LSPs	Yes	No	Yes for PCEP attributes	Yes	Yes for non-PCEP attributes such as design flags	No
PCE-initiated LSPs	Yes	No	Yes for PCEP attributes	Yes	Yes for non-PCEP attributes such as design flags	No

Table 53: Effects of Resetting or Synchronizing the Network Model (Continued)

Element	Is the element removed from the database?		Is the item sent back to the controller by the live network?		Could data be lost?	
	Reset	Sync	Reset	Sync	Reset	Sync
Multilayer nodes	Yes	No	Yes	No	Yes for designed attributes such as user-defined names	No
Multilayer links	Yes	No	Yes	No	Yes for design attributes such as Comment	No
Interlayer links	Yes	No	No	Yes, links mapped to known nodes are re-sent.	Yes	Yes, access links to unknown nodes are lost and need to be recreated
Multilayer-derived facilities	Yes	No	Yes	No	No	No
Link-derived facilities	Yes	Yes	Yes	Yes	Yes	Yes
Ongoing maintenance events	No	No	N/A	N/A	No	No

Table 53: Effects of Resetting or Synchronizing the Network Model (Continued)

Element	Is the element removed from the database?		Is the item sent back to the controller by the live network?		Could data be lost?	
	Reset	Sync	Reset	Sync	Reset	Sync
Future maintenance events	Yes	No	N/A	N/A	Yes	No
Ongoing scheduled LSPs	No	No	N/A	N/A	Yes (scheduled LSP is never terminated)	No
Future scheduled LSPs	Yes	No	N/A	N/A	Yes	No
Device profiles	No	No	N/A	N/A	No	No
Router latitude and longitude	No	No	N/A	N/A	No	No
Router grouping	No	No	N/A	N/A	No	No
Users table	No	No	N/A	N/A	No	No
Saved map layout	No	No	N/A	N/A	No	No
Events	No	No	N/A	N/A	No	No
Scheduled path optimization	No	No	N/A	N/A	No	No



WARNING: Perform this action only under the supervision of JTAC. This action erases the network data model and the data provided through the Add or Modify actions in the network information table (user model data).

RELATED DOCUMENTATION

[Paragon Pathfinder Overview](#) | 6

Disaster Recovery Overview

IN THIS SECTION

- [Failure Scenarios](#) | 255

Disaster recovery is the deployment of Paragon Automation cluster at two different geographical locations. Disaster recovery ensures that when the Paragon Pathfinder component in one Paragon Automation deployment goes down, Paragon Pathfinder services are available from the other Paragon Pathfinder component in the other Paragon Automation deployment.

In a disaster recovery setup, you must independently configure network discovery (for example, BGP-LS, PCEP, and analytic streaming) and independently perform tasks such as discovering devices and configuring playbooks, on both the Paragon Automation deployments.

A federated exchange configured on the messaging bus synchronizes the two deployments for:

- Topology-related changes (changes related to addition, modification and deletion of nodes, links, and LSPs)
- LSP optimization
- Resetting the topology
- Adding or removing LSP delegation

The toposervers in the deployments send beacons over a federated exchange which help in detecting failure of either of the deployments and avoid unsynchronized changes. By default, the toposervers send beacons once every 5 seconds.

NOTE:

- You can provision or modify only path computation client (PCC)-delegated LSPs in a disaster recovery setup. Provisioning or modification of PCC-controlled LSPs and path computation element (PCE)-initiated LSPs is not supported.
- NETCONF is not supported in a disaster recovery setup.
- You must run device collection task for delegated LSPs to appear correctly on both the deployments. Otherwise, the LSP appears as PCC-controlled in the deployment without the delegation bit and modifying the LSP from the deployment without the delegation bit results in an error.

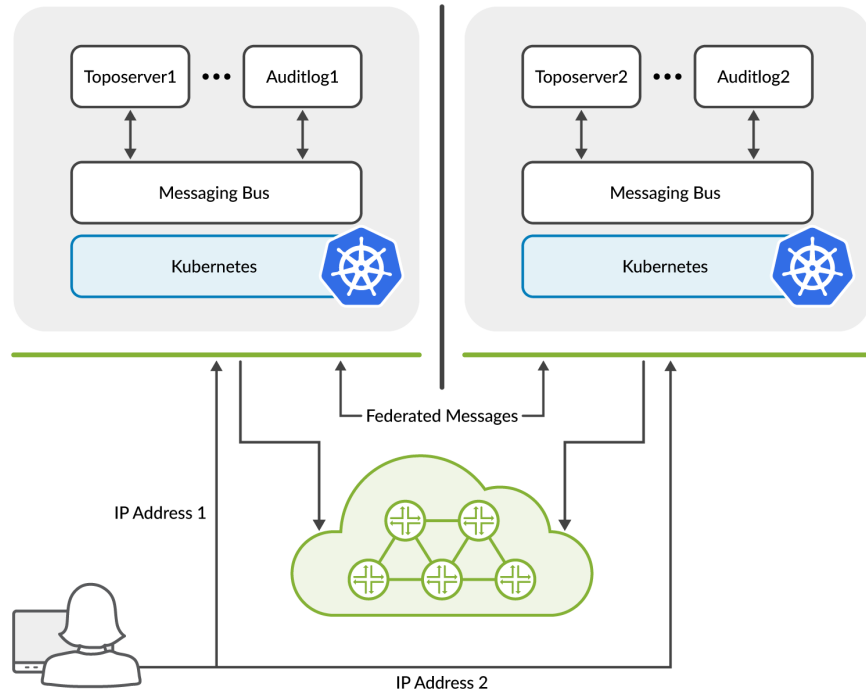
For a disaster recovery setup, you must configure the following flags in the Pathfinder Settings (**Configuration > Network Settings > Pathfinder Settings**):

- source-of-truth (optional)
- use-federated-exchange (mandatory)
- toposerver-beacon-interval (optional)
- sync-topology-after-failure (optional)

For information on configuring the Pathfinder settings, see ["Modify Pathfinder Settings From the GUI" on page 206](#).

[Fig on page 254](#) shows the disaster recovery architecture, where Paragon Pathfinder is deployed at two different geographical locations to manage the same network. Both the sites have active BGP Link State (BGP-LS), Path Computation Element Protocol (PCEP), BGP Monitoring Protocol (BMP), and SSH sessions to the network. The deployments communicate with the network through a federated exchange. Both the deployments are in active-active state and synchronize changes within a few seconds. For information on configuring disaster recovery, see [Configure Disaster Recovery for Paragon Pathfinder](#).


Figure 9: Disaster Recovery Architecture



In a disaster recovery setup, we recommend that you specify one of the deployments as the primary deployment by configuring the delegation priority for the path computation element (PCE). All delegated LSPs send their delegation bit to the PCE in the primary deployment. To configure delegation priority for a PCE, see [delegation-priority](#).

[Figure on page 255](#) shows the PCEP connectivity between the Paragon Pathfinder deployments (clusters) at two different locations and the network. The LSPs are delegated to either cluster 1 or cluster 2. Notifications related to the LSPs are reported to both the clusters.

Figure 10: PCEP Connectivity in a Disaster Recovery Setup



Failure Scenarios

The following failure scenarios are possible:

- The primary deployment fails and PCC sends delegation bit to the PCE in the secondary deployment.

When the primary deployment fails, the toposerver in the secondary deployment goes into safe mode preventing input of any change related to topology, This ensures that the topology information in the deployments do not go out of synchronization.

If you want to make changes while the primary deployment is down and not sending beacon, enable the source-of-truth flag in Paragon Pathfinder settings (**Configuration > Network Settings > Pathfinder Settings**) on the secondary deployment. See ["Modify Pathfinder Settings From the GUI" on page 206](#) for details. After you bring the primary deployment up, disable the source-of-truth flag on the primary deployment (if enabled) so that the primary deployment synchronizes with the secondary deployment to obtain all changes that were made through the secondary deployment. After the synchronization, we recommend that you disable the source-of-truth flag on the secondary deployment.

- Federated exchange fails.

The federated exchange transfers messages between the individual deployments. If the federated exchange fails, Paragon Pathfinder cannot execute provisioning and modification requests. To check if the federated exchange is up or not, execute the following command on any deployment:

```
root@davinci-master-c1:~# kubectl exec -it -n northstar rabbitmq-0 -- rabbitmqctl
list_parameters
```

An output in the below format indicates that the federated exchange is up.

```
Listing runtime parameters for vhost "/" ...
component      name      value
federation-upstream  my-upstream  {"expires":30000,"uri":"amqps://
northstar:6rchz9eHGy@10.52.33.97?cacertfile=/opt/bitnami/rabbitmq/certs/
ca_certificate.pem&verify=verify_none"}
```

If the federated exchange fails, you must troubleshoot and bring the federated exchange up to manage the network.

After you bring the federated exchange up, the toposervers at the primary deployment and the secondary deployment synchronize any changes to the topology.

RELATED DOCUMENTATION

| [Understand LSP Delegation and Undelegation](#) | 815

Network Slicing Overview


IN THIS SECTION

- [Benefits of Network Slicing](#) | 259

Network slicing enables network operators to define logical networks on a physical network. The logical networks are called network slices. The slices have a shared control plane and can have a dedicated data plane. You can define shared logical slices as an additional administrative constraint to steer LSPs.

A slice comprises a set of nodes, links, and prefixes of a transport network; see [Figure 11 on page 257](#). Each slice is managed by a single network operator.

Figure 11: Example of Slices in a Network




You can configure network slices by using the GUI or REST APIs. The nodes and the links belonging to a slice are grouped into a topology group, and the slice (topology group) is identified by a group ID (also referred to as slice ID). A node and a link can be part of more than one slice (upto thousand slices). For information about configuring and working with a slice, see ["Group Nodes and Links into a Topology Group" on page 683](#).

Label-switched paths (LSPs) are routed through nodes and links that have the same slice IDs. To constrain an LSP to traverse through a particular slice, you must assign the slice ID of that slice to the LSP. LSPs that do not have an assigned slice ID can be routed through any node or link, irrespective of whether a node or link is part of a slice. When a node or link in a slice is down, the LSPs are rerouted through a redundant path in the same slice as shown in [Figure 12 on page 258](#) and [Figure 13 on page 259](#).


In [Figure 12 on page 258](#), nodes vmx101 and vmx103 are connected through the slice with slice ID 827. The slice includes nodes vmx101, vmx102, vmx104, vmx105, vmx106, ios-xr8, and ios-xr9, and the links between these nodes. The traffic flows between vmx101 to vmx103 through a path comprising nodes vmx105, vmx102, vmx106, vmx104, and vmx107.

Figure 12: LSP Routed through a Slice



In [Figure 13 on page 259](#), the link between vmx102 and vmx106 is down, as indicated by the red F on the link. The traffic between vmx101 and vmx103 now flows through a redundant path within slice 827, which comprises nodes vmx105, vmx102, ios-xr8, ios-xr9, vmx104, and vmx107.

Figure 13: LSP Rerouted through the Same Slice When a Link is Down



Benefits of Network Slicing

- Network slicing enables network operators to deliver services that have competing service requirements over a shared infrastructure.
- As network slicing allows differentiation of performance characteristics for different types of slices, the applications which need the best performance gets the best performance..

RELATED DOCUMENTATION

[About the Topology Group Tab | 801](#)


[Group Nodes and Links into a Topology Group | 683](#)

Add a Test Agent for Network Slices

While provisioning a network slice, Paragon Active Assurance (PAA) runs tests on the slice by using Test Agents. The Test Agents are connected to the nodes before configuring a network slice and the connections are managed in Paragon Pathfinder topology. The Test Agent generates active, synthetic traffic for measuring service quality in the slice across multiple applications, services, and technology domains. For information about Paragon Active Assurance, see [Paragon Active Assurance User Guide](#).

Figure on page 260 shows a topology with Test Agents PAA-Test-1 and PAA-Test-2 connected to node vmx103 and Test Agent PAA-Test-3 connected to node vmx101.

Figure 14: Network Topology with Paragon Active Assurance Test Agents



To add a Test Agent to a network:

1. Add a Test Agent to the network topology; see ["Add a Node" on page 696](#).

NOTE: To add a Test Agent, select Active Assurance for the node's role.

2. Add a link from the Test Agent to a provider edge node in the network topology; see ["Add a Link" on page 705](#).

RELATED DOCUMENTATION

[Network Slicing Overview](#) | 256

Configure LSP Routing in a Network Slice by Using a Path Computation Profile

A path computation profile defines the set of traffic engineering (TE) constraints (for example, admin color, cost, and delay) that Paragon Pathfinder uses to compute label-switched-paths (LSPs).

NOTE: To configure LSP routing in a network slice by using computation profiles, you must have prior knowledge about:

- General TE concepts
- RFC7950, *The YANG 1.1 Data Modeling Language*
- RFC7951, *JSON Encoding of Data Modeled with YANG*
- RFC8040, *Internet Web Replication and Caching Taxonomy*
- RESTCONF protocol

You can use path computation profiles for deciding the routing of LSPs within a network slice. For example, you can define a profile for routing LSPs with admin color 150 through a slice with the slice ID 100 or, you can also define a profile to route LSPs with slice ID 100 with a delay of 10 ms.

NOTE: You can route PCC-delegated and PCC-controlled LSPs by using path computation profiles. You cannot route RSVP LSPs and PCE-initiated LSPs by using computational profiles.

Before you create a profile, ensure that you have access to and understand the following:

- juniper-pathfinder-lsp-policy.yang
- juniper-pathfinder-profile.yang
- lsp.json (REST LSP model)

To use a path computation profile for configuring network slices, you must:

1. Create a profile (defined by the juniper-pathfinder-profile YANG model) in the RESTCONF interface .
2. Configure a policy (defined by the juniper-pathfinder-lsp-policy YANG model) to map the LSP to the profile by using the RESTCONF interface .

The following example profile routes LSPs on links that have lower delay.

```
{
  "juniper-pathfinder-profile:computation-profiles": [
    {
      "comment": "Profile low-delay in use",
      "id": "low-delay",
      "path-affinities-values": {
        "path-affinities-value": [
          {
            "usage": "resource-aff-include-any",
            "value": "02"
          }
        ]
      }
    }
  ]
}
```

The following sample of a policy uses the low-delay computation profile on LSPs that have

- String DEMO in their name, or
- SR-TE policy defined with admin color or SR-TE color between 100 and 200

You can interpret the policy condition by using [grule](#). The LSP data follows the Paragon Pathfinder REST API model.

```
{
  "juniper-pathfinder-lsp-policy:lsp-policy": {
    "policy-definitions": {
      "policy-definition": [{
        "actions": {
          "path-computation-profiles": [
            "low-delay",
            "prio-5"
          ]
        },
        "comment": "Use low delay for some LSPs(testMe)",
        "conditions": {
          "condition": "( (LSP.liveProperties != nil) && (LSP.liveProperties.srPolicy != nil) && ((LSP.liveProperties.srPolicy.color >= 100 || LSP.liveProperties.srPolicy.color <= 200 ))) || LSP.name.Contains(\"DEMO\")"
        }
      ]
    }
  }
}
```

```
    },  
    "name": "UseLowDelay",  
    "priority": 20,  
    "terminal": true  
  }]  
}  
}
```

RELATED DOCUMENTATION

[Network Slicing Overview](#) | 256

Network Groups

IN THIS CHAPTER

- [About the Network Groups Page | 264](#)
- [Add a Network Group | 266](#)
- [Edit a Network Group | 269](#)

About the Network Groups Page

IN THIS SECTION

- [Tasks You Can Perform | 265](#)

To access network group page in Paragon Automation Platform GUI, click **Configuration > Network Groups**.

Network groups allow you to correlate health status data between multiple devices across the network. For example, you can create a network group that monitors the ping times between two or more devices and notifies you if the ping times are too high.

[Table 54 on page 264](#) describes the fields in the Network Groups page.

Table 54: Fields in Network Groups Page

Fields	Description
Network Name	Displays the name of the network group.

Table 54: Fields in Network Groups Page (Continued)

Fields	Description
Description	Displays the description you give when adding the network group.
Playbooks	Displays the number of playbooks applied to the network group.
Instances	Displays the number of playbook instances applied to the network group.
Logging	<p>Displays the first letter of the log severity such as Debug, Error, Warn or Info along with the log level.</p> <p>The log level for the device group shows Global if you configure device group to collect logs for every service running on a device group.</p> <p>If you configure service specific logs, the column displays Others.</p>

Tasks You Can Perform

You can perform the following tasks in this page:

- Add a network group. See ["Add a Network Group" on page 266](#).
- Edit a network group. See ["Edit a Network Group" on page 269](#).
- Export details of all network groups or of a particular network group.

To export configuration details of all network groups:

1. Go to **Configuration > Network Groups**.

You are taken to the Network Configuration page.

2. Click on the **Export** button (at the top right corner) and select *Export as CSV* from the menu.

In the pop up window, click **Open** to view the Excel file or **Save As** to save the Excel file to any location in your system.

RELATED DOCUMENTATION

| [Enable Alert Notifications for Device Groups and Network Groups](#) | 627

Add a Network Group

To add a network group:

1. Click the **Configuration > Network** option in the left-navigation bar.
2. Click the + (Add Network) button.
3. Enter the necessary values in the text boxes and select the appropriate options for the network group.

The following table describes the attributes in the **Add a Network Group** window:

Table 55:

Attributes	Description
Name	Name of the network group. (Required)
Description	Description of the network group.
Reports	<p>In the Reports field, select one or more health report profile names from the drop-down list to generate reports for the network group. Reports include alarm statistics, device health data, as well as device-specific information (such as hardware and software specifications).</p> <p>To edit or view details about saved health report profiles, go to the Ingest Settings page under the Administration menu in the left navigation bar. The report profiles are listed under Report Settings page (Report tab).</p> <p>For more information, see "Configure Report Settings" on page 628.</p>
Disable Trigger Action Scheduler (for a particular network group)	<p>By default, this field is marked False because the option to add a UDA scheduler in Trigger Action page is enabled.</p> <p>You can set this field to True if you want to disable UDA scheduler settings in Trigger Action page.</p>

Table 55: (Continued)

Attributes	Description
Notifications	<ul style="list-style-type: none"> You can use the Alert page to organize, track, and manage KPI alarm notifications received from Paragon Insights devices. To receive Paragon Insights alert notifications for KPI events that have occurred on your devices, you must first configure the notification delivery method for each KPI event severity level (Major, Minor, and Normal). Select the delivery method from the drop-down lists. <p>To edit or view details about saved delivery method profiles, go to the Ingest Settings page under the Administration menu in the left navigation bar. The delivery method profiles are listed under Notification Settings (Notification tab).</p>
Ingest Frequency	Select existing Ingest Frequency Profiles to override rule or sensor frequency settings.
Logging Configuration	<p>You can collect different severity levels of logs for the running Paragon Insights services of a network group. Use these fields to configure which log levels to collect:</p> <p>Global Log Level From the drop-down list, select the level of the log messages that you want to collect for every running Paragon Insights service for the network group. The level is set to error by default.</p> <p>Log Level for specific services Select the log level from the drop-down list for any specific service that you want to configure differently from the Global Log Level setting. The log level that you select for a specific service takes precedence over the Global Log Level setting.</p>

Table 55: (Continued)

Attributes	Description
Publish	<p>You can configure Paragon Insights to publish Paragon Insights sensor and field data for a specific network group:</p> <p>Destinations Select the publishing profiles that define the notification type requirements (such as authentication parameters) for publishing the data.</p> <p>To edit or view details about saved publishing profiles, go to the System page under the Settings menu in the left-nav bar. The publishing profiles are listed under Notification Settings.</p> <p>Field Select the Paragon Insights rule topic and rule name pairs that contain the field data you want to publish.</p>
Tagging	<p>Select one or more tagging profiles from the existing profiles list. Tagging makes use of profiles to set conditions, define new fields and keys, and insert values into those fields after creation. For more information, see "Paragon Insights Tagging Overview" on page 570.</p>
Root Cause Analysis	<p>Mode is enabled by default. Disable Mode if you do not want the device group to be a part of resource discovery and dependency process.</p> <p>Exclude Resources field allows you to select network resources that must be excluded from resource and dependency formation.</p>

4. Click **Save** to save the configuration or click **Save and Deploy** to save and deploy the configuration.

RELATED DOCUMENTATION

| [Edit a Network Group](#) | 269

Edit a Network Group

To edit a network group:

1. Click the **Configuration > Network Groups** option in the left navigation bar.
2. Click anywhere on the line that contains the group name in the table under **NETWORK LIST**.
3. Click on the **Edit Network (Pencil)** icon.
4. Modify the attributes, as needed.
See "[Add a Network Group](#)" on page 266 for a description of each attribute.
5. Click **Save** to save the configuration or click **Save and Deploy** to save and deploy the configuration. .
6. (Optional) A network can be deleted by clicking the **Delete Network (Trash Can)** icon.

RELATED DOCUMENTATION

| [About the Network Groups Page](#) | 264

Topology Filter

IN THIS CHAPTER

- [About the Topology Filter Page | 270](#)
- [Add a Topology Filter Rule | 274](#)
- [Edit a Topology Filter Rule | 275](#)
- [Delete a Topology Filter Rule | 276](#)

About the Topology Filter Page


IN THIS SECTION

- [Enable the Topology Filter Service | 271](#)
- [Tasks You Can Perform | 272](#)
- [Field Descriptions | 273](#)

To access this page, click **Configuration > Network Settings > Topology Filter**.

The topology filter service allows you to limit the nodes appearing in your topology map to a subset of the nodes in your network. This capability is important when your network contains more nodes than your license covers and you want to control which nodes are recognized. You might also want to filter out nodes that are not important for traffic engineering management such as aggregation layer nodes or route reflectors.

The topology filter is available only in installations utilizing BGP Monitoring Protocol (BMP) for the topology acquisition method. The topology filter is added to the BMP messaging pipeline between the BMP monitor and the topology server.



Filtering the topology involves creating a set of rules. Each rule consists of the field (Condition Field) to search on, the value (Condition Value) to look for, the Condition Type to be applied (match or regular expression) and the Action to be taken if the value is a match. The rules are automatically applied in a sequence to the topology map. Once a node is matched based on the rule, it cannot be matched by a subsequent rule. The highest rule or the rule at the top of the list of topology filter rules, controls this decision. For example, once a decision is made about the node (accept or reject), it will not change even if a rule below in the list provides the opposite result. The topology filter service checks links and removes links that involve an rejected node, even if the rejected node is reporting information to the traffic engineering database.

Enable the Topology Filter Service

Before using the topology filter in the UI, you must first enable the topology filter service by modifying the configuration file.

You must modify the BMP host and port values in the `northstar-config` configuration map file using a text editing tool such as `vi`.

To enable the topology filter service:

1. Log in to the Paragon Automation primary node.
2. Edit the `northstar-config` file by executing

```
kubectl edit cm -n northstar northstar-config
```

3. Set the following parameters:

- `bmp_host=ns-filter`

(Default is `bmp-grpc` which is the `bmpMonitor` endpoint.)

- `bmp_port=10004`

(Default is 10002.)

4. Save the `northstar-config` file and exit.

All the pods using this configuration file are restarted automatically.

NOTE: You can confirm if the status of the services are up by executing:

```
kubectl -n northstar get pods
```

Tasks You Can Perform

You can perform the following tasks from this page:

- View the existing the topology filters rules. For more details, click **More** > **Detail** or click the **Details** icon. The **Result Of Applied Filters** page appears.
- Add a topology filter rule. See ["Add a Topology Filter Rule" on page 274](#).
- Edit a topology filter rule. See ["Edit a Topology Filter Rule" on page 275](#).
- Delete a topology filter rule. See ["Delete a Topology Filter Rule" on page 276](#).
- View detailed information about the topology filter rule by clicking the details icon that appears when you hover over the filter rule or select **More** > **View Details**. A **Result of Applied Filters** page appears on the right displaying the filter rule details.
- Export all the topology filter rules to a CSV file by clicking **More** > **Export Rules**.
- Import topology filter rules from a CSV file by clicking **More** > **Import Rules**. On the **Import Filter Rule** page, click **Browse** and select the CSV file that you want to import. Click **OK**.

The **Filter Rules** table is updated with the new topology filter rules from the CSV file.

NOTE: If you have multiple rules to add, we recommend that you:

1. Export the current topology rules to a CSV file.
2. Update that CSV file with the new rules.

3. Import the updated CSV file.

- Show/Hide Columns: Choose to show or hide a specific column in the Filter Rules table.

Hover over the **More Options** (vertical ellipses) > **Show/Hide Columns** and select the *Column-Name* check box of the column you want to display in the table.

Field Descriptions

Table 56: Fields on Topology Filter Page

Field	Description
Condition	<p>Select one of the following fields as a condition for filtering a node:</p> <ul style="list-style-type: none"> • hostname • router-id • sys-id
Type	Type of value to be matched. You can select either text or regular expression to match the condition.
Value	Value to be matched.
Action	<p>Action to be performed when the condition value matches the condition set. You can select either Accept or Reject.</p> <p>If you select:</p> <ul style="list-style-type: none"> • Accept—Nodes matching the condition are displayed on the topology map. • Reject—Nodes not matching the condition are not displayed on the topology map.

RELATED DOCUMENTATION

[About the Topology Page | 677](#)

[Interactive Map Features Overview | 660](#)

Add a Topology Filter Rule

Before you add a topology filter rule, ensure that you have enabled the topology filter service. See ["Enable the Topology Filter Service" on page 271](#).

To add a topology filter rule:

1. Click **Configuration > Network Settings > Topology Filter**.

The Topology Filter page appears.

2. Click the add (+) icon.

The Add New Filter page appears.

3. Configure the fields as per [Table 57 on page 274](#).

4. Click **OK**.

The new rule is listed on the Topology Filter page.

NOTE: To change the order of the rules, you can drag and drop the rows to shift them (up or down). The rules sequence or order of rules is significant. The rule at the top of the list has the highest priority when making decisions for the nodes.

5. Click **Save Changes** to save the rules and apply the filter rule to the topology map.

A confirmation message appears indicating that the topology filter data is saved.

6. Navigate to the Topology (**Network > Topology**) page to view the changes in the topology map based on the filter rule you applied.

You can view the filtered status on the **Result of Applied Filters** page by clicking **View > Details** or clicking the details icon that appears when you hover over the filter row. Every node in the full topology is listed. The **Filtered** column indicates **Yes** if the node was filtered out (excluded), and **No** if the node was not filtered (still included in the topology).

Table 57: Fields on the Add New Filter Page

Field	Description
Condition Field	Select one of the following fields as a condition for filtering a node: <ul style="list-style-type: none"> • hostname • router-id • sys-id

Table 57: Fields on the Add New Filter Page (Continued)

Field	Description
Condition Type	Select the type of condition; match or regex (regular expression). For router-id, only match is allowed.
Condition Value	Enter the value to be matched.
Action	Select an Action; either Accept or Reject to include or exclude the node from the topology respectively.

RELATED DOCUMENTATION

| [About the Topology Page | 677](#)

Edit a Topology Filter Rule

To edit a topology filter rule:

1. Click **Configuration > Network Settings > Topology Filter**.
The Topology Filter page appears.
2. Select the topology filter rule that you want to edit.
3. Click the edit (**pencil**) icon.
The Edit Filter Rule page appears.
4. Configure the fields as per [Fields on the Add New Filter Page on page 274](#).

NOTE: If you select multiple rules for modification, you can edit only the **Action** field (Accept or Reject). If you select a single rule, you can edit all the fields.

5. Click **OK**.
The edited topology filter rule is displayed on the Topology Filter page.
6. Click **Save Changes** to save the rules and apply the filter rule to the topology map.

Navigate to the Topology (**Network > Topology**) page to view the changes in the topology map based on the filter rule you applied.

RELATED DOCUMENTATION

[About the Topology Filter Page](#) | 270

Delete a Topology Filter Rule

To delete a topology filter rule:

1. Click **Configuration > Network Settings > Topology Filter**.

The Topology Filter page appears.

2. Select the topology filter rule that you want to delete.

NOTE: You can delete multiple filter rules at a time.

3. Click the delete (**trash can**) icon.

A confirmation dialog box appears.

4. Click **Yes**.

The topology filter rule is removed from the Topology Filter page.

5

PART

Manage Device Templates and Configuration Templates

[Configuration Templates](#) | 278

[Device Templates](#) | 298

Configuration Templates

IN THIS CHAPTER

- [Configuration Templates Overview | 278](#)
- [Configuration Templates Workflow | 280](#)
- [About the Configuration Templates Page | 281](#)
- [Add Configuration Templates | 284](#)
- [Preview and Render a Configuration Template | 291](#)
- [Assign Configuration Templates to a Device Template | 292](#)
- [Deploy a Configuration Template to a Device | 293](#)
- [Edit, Clone, and Delete a Configuration Template | 295](#)

Configuration Templates Overview

IN THIS SECTION

- [Benefits | 280](#)

Paragon Automation provides configuration templates to provision opaque configurations, both during onboarding and throughout the device lifecycle, for Juniper Networks and other third-party devices. By using configuration templates, you can deploy customized configurations on devices that are managed by Paragon Automation.

A configuration template can be used as follows:

- **Globally**—You can define the configuration (for example, SNMP Configuration) to be applied to all the devices managed by Paragon Automation.

- Device-specific—You can define a configuration that is specific to a device; for example, BGP configuration

By default, Paragon Automation provides some predefined configuration templates. See [Table 58 on page 279](#) for the list of the predefined configuration templates. You can also create your own templates by cloning an existing template and modifying its settings. Templates can be added by administrators or users with privilege to add configuration templates. Some of the predefined templates are pre-assigned to specific device templates for enabling configurations that are required during onboarding of a device; for example, username and password.

[Table 58 on page 279](#) lists the predefined configuration templates available in Paragon Automation.

Table 58: Predefined Configuration Templates in Paragon Automation

Name	Description
AE_DEVICE_COUNT	Configure the aggregated Ethernet interfaces on a device.
BANNER	Configure the banner that appears when you log in to a device.
DNS	Configure Domain Name System (DNS) server settings on a device.
DOMAIN_NAME	Configure the domain name on a device.
HOSTNAME	Configure the host name on a device.
LLDP	Enable and configure Link Layer Discovery Protocol (LLDP) on all interfaces of a device.
LOCAL_USER	Configure a local user on a device.
NETCONF	Configure NETCONF on a device.
NTP	Configure Network Time Protocol (NTP) settings on a device.
SNMP	Configure basic SNMP version 2 (SNMPv2) parameters on a device.
SSH	Configure SSH parameters on a device.

Table 58: Predefined Configuration Templates in Paragon Automation (Continued)

Name	Description
SYSLOG	Configure system log settings on a device.

You can deploy a configuration template as follows:

- Deploy the configuration template directly on a device; see ["Deploy a Configuration Template to a Device" on page 293](#)
- Assign the configuration template to a device template so that the configuration is deployed on a device during zero-touch provisioning (ZTP) and device discovery; see ["Assign Configuration Templates to a Device Template" on page 292](#).

Benefits

Configuration templates provide a mechanism to create customized configurations and push the configurations to one or more devices. This helps you to deploy configurations beyond the standard configuration templates provided in Paragon Automation.

RELATED DOCUMENTATION

[Devices Overview | 130](#)

[Device Templates Overview | 298](#)

[View and Manage Device Configuration | 165](#)

Configuration Templates Workflow

The high-level workflow for configuration templates is as follows:

1. You can use a pre-existing template (skip to step 2) or create a new template using one of the following methods:
 - Clone an existing configuration template and modify the cloned template. For more information, see ["Edit, Clone, and Delete a Configuration Template" on page 295](#).
 - Add a configuration template by specifying the template configuration and logic. For more information, see ["Add Configuration Templates" on page 284](#).

2. (Optional) We recommend that you preview and validate the configuration template before assigning the configuration template to a device template or deploying the configuration template directly on a device. For more information, see "[Preview and Render a Configuration Template](#)" on page 291.
3. You can assign a configuration template to a device template from the Device Templates page. This enables you to deploy additional configuration on the device during zero touch provisioning (ZTP) or during device discovery, and after the device is onboarded. For more information, see "[Deploy a Configuration Template to a Device](#)" on page 293.
4. You can deploy a configuration template directly on one or more devices that were previously activated, which enables you to deploy templates that were added after a device was activated or to deploy additional configuration to devices. You can deploy configuration templates to devices from the Configuration Templates page. For more information, see "[Assign Configuration Templates to a Device Template](#)" on page 292.

RELATED DOCUMENTATION

[Edit Configuration Templates Assigned to a Device Template](#) | 302

[View and Manage Device Configuration](#) | 165

About the Configuration Templates Page

IN THIS SECTION

- [Tasks You Can Perform](#) | 281
- [Field Descriptions](#) | 283

To access this page, click **Configuration > Templates > Config Templates** in the left navigation menu.

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 *<TemplateName>* pane appears on the right side of the page displaying the configuration template details.

- Add a configuration template; see ["Add Configuration Templates" on page 284](#).
- Preview a configuration template; see ["Preview and Render a Configuration Template" on page 291](#).
- Deploy a configuration template on one or more devices; see ["Deploy a Configuration Template to a Device" on page 293](#).
- Assign a configuration template to a device template; see ["Assign Configuration Templates to a Device Template" on page 292](#).
- Clone, edit, or delete a configuration template; see ["Edit, Clone, and Delete a Configuration Template" on page 295](#).
- Filter Entries—Filter the table entries by adding new filtering criteria.

Hover over the **Filter** (funnel) icon and select **Add Filter**. On the **Add Criteria** page, select the filtering criteria from the **Field** and **Condition** list, and enter the text to be compared in the **Value** field. Click **Add**.

The filtered table entries are listed and the filter criteria name is displayed above the table column names.

You can click the cross (X) icon (next to the filter name) to remove the filtering criteria.

NOTE:

- You can add multiple filtering criteria. Once you add the multiple filtering criteria, you can select the **And** condition to display the entries matching all the filtering criteria or select the **Or** condition to display the entries matching any one of the filtering criteria.
- Quick filter: Save the filtering criteria as quick filters. Once you have added all the filtering criteria, you can save a particular criteria or multiple criteria for future use by clicking **Save**.

On the **Save Filter** window, 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 a specific column in the table.
Hover over the **More Options** (vertical ellipsis) > **Show/Hide Columns** and select the *Column-Name* check box of the columns you want to display in the table.
- Reset Preference—Reset the displayed columns to the default set of columns for each tab in the table.

Hover over the **More Options** (vertical ellipsis) and select **Reset Preference**.

- **Sort Entries**—Click the column name to highlight the up and down arrows next to the column name. Sort the table entries in ascending or descending order of that column by clicking the up or down arrow respectively.

Field Descriptions

[Table 59 on page 283](#) displays the fields on the Configuration Templates page.

Table 59: Fields on the Configuration Templates Page

Field	Description
Name	The name of the configuration template.
Format	The format in which the configuration template is defined—CLI or XML.
Family	The device family for which the configuration template is applicable: <ul style="list-style-type: none"> • Cisco IOS-XR • Juniper-ACX • Juniper-EX • Juniper-MX • Juniper-PTX • Juniper-QFX • Juniper-SRX • Juniper-ANY
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 format.

Table 59: Fields on the Configuration Templates Page (*Continued*)

Field	Description
Created by	The user who created the configuration template. <i>System</i> indicates that the template is a predefined template.

RELATED DOCUMENTATION

[Devices Overview | 130](#)

[Device Templates Overview | 298](#)

Add Configuration Templates

To add a configuration template, you should either be a user with administrative privileges or have the privilege to add configuration templates.

NOTE:

- Before you add the configuration template, ensure that you have the device configuration ready.
- We recommend that you use a working device configuration to add the configuration template.

To add a configuration template:

1. Select **Configuration > Templates > Config Templates** on the left navigation menu.
The Configuration Templates page appears.
2. Click the **Add** icon (+).
The Add Configuration Template page (wizard) 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 60 on page 285](#).

4. Click **Next** to go to the Template Configuration tab.

5. Add the configuration on the Template Configuration tab.

You can view a sample configuration by clicking the **Sample Configuration** link.

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.
- Refer to the sample configuration file for adding the configuration.
- Parameterize variables by using double curly braces `{{}}`.

6. Click **Next** to go to the Generated UI tab, where you can view the UI for the parameters that you entered.

7. Perform one or more actions on the Generated UI tab, as explained in [Table 61 on page 286](#).

8. Click **Save**.

The configuration template is added and you are returned to the Configuration Templates page, where a confirmation message appears.

[Table 60 on page 285](#) lists fields to be entered on the Basic Information tab of the Add Configuration Templates page.

Table 60: 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 and hyphens; 64-characters maximum.
Description	Enter a description for the configuration template; 255-characters maximum.
Configuration Format	Select the output format for the configuration template: <ul style="list-style-type: none"> • CLI (default) • XML

Table 60: Fields on the Basic Information Tab of the Add Configuration Templates Page (Continued)

Field	Description
Device Family	<p>Select a device family for which you are adding the template.</p> <ul style="list-style-type: none"> • Cisco-IOS-XR • Juniper-ACX • Juniper-ANY (the configuration can be pushed to any Juniper Networks' devices) • Juniper-EX • Juniper-MX • Juniper-PTX • Juniper-QFX • Juniper-SRX • NOKIA

[Table 61 on page 286](#) lists the actions that you can perform on the Generated UI tab of the Add Configuration Templates page.

Table 61: 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 61: Generated UI Actions (Add Configuration Template Page) (Continued)

Action	Description
Modify the settings for a field, section, or grid	<p>To modify the settings for a field, section, or grid:</p> <ol style="list-style-type: none"> 1. Click the Settings (gear) icon next to the field, section, or grid. <p>The Form Settings pane appears on the right side of the page, displaying the Basic Settings and Advanced Settings tabs.</p> <ol style="list-style-type: none"> 2. Modify the fields on these tabs, as needed. See Table 62 on page 288 for an explanation of the fields on these tabs. 3. Click Save Settings for each field to save your changes. <p>The modifications that you made are displayed on the UI.</p>
Reset the generated UI	<p>Click Undo all Edits to discard the changes that you made and undo the changes made on the UI.</p>
Preview configuration	<p>Preview the configuration defined in the configuration template.</p> <p>To preview a configuration template:</p> <ol style="list-style-type: none"> 1. Click Preview Configuration. <p>The Preview Configuration page appears, displaying the configuration that was rendered based on the values that you entered.</p> <ol style="list-style-type: none"> 2. Check if the configuration was rendered correctly. <ul style="list-style-type: none"> • If the configuration was not rendered correctly, click the close (X) icon to go back and make modifications as needed. • If the configuration was rendered correctly, click OK. <p>You are returned to the Generated UI page.</p>

[Table 62 on page 288](#) lists the fields on the Form Settings pane.

Table 62: 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.
Input Type	<p>Select the input type for the parameter in the configuration template:</p> <ul style="list-style-type: none"> • 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 (Boolean): If the input value for the parameter is a boolean value (true or false). • 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 62: Form Settings (Add Configuration Template Page) (Continued)

Setting	Guideline
Validate	<p>For Text input type, select one or more validation criteria against which the input value will be checked:</p> <ul style="list-style-type: none"> • No Space • Alpha and Numeric • Alpha, Numeric, and Dash • Alpha, Numeric, and Underscore <p>If the value that you entered for the parameter on the UI does not meet the selected validation criteria, an error message appears.</p> <p>NOTE: For greater control of input values, you can use the regular expression option in the Advanced Settings tab.</p>
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	Click the toggle button to make the parameter common across all devices to which the configuration template is being deployed. If you disable the toggle button, which is default, the parameter must be specified for each device.
Hidden	<p>Click the toggle button to hide the parameter on the UI when you preview and deploy the template.</p> <p>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.</p>
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 62: 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).
Resource Type	<p>For Dropdown input type, select the type of resource from which you want to retrieve data:</p> <ul style="list-style-type: none"> • Static Resource—Resources in the list on the UI are mapped to the values that you specify. <ul style="list-style-type: none"> • To add a static resource: <ol style="list-style-type: none"> 1. Click the Add (+) icon. Cells appear in the List Values table. 2. Click inside the cells to specify values for the Label (name for the option in the list), Value (value for the option in the list), and Visibility (conditional visibility based on the option selected from the list) fields. 3. Click ✓ (check mark) to save your changes. The values that you specified are displayed in the List Values table. • To edit a static resource, select the resource and click the Edit (pencil) icon. • To delete a static resource, select the resource and click the Delete (X) icon.

Table 62: Form Settings (Add Configuration Template Page) (Continued)

Advanced Settings Tab

Regexp	<p>Enter a regular expression (regex pattern) to validate the input value.</p> <p>A regular expression defines a search pattern that is used to match characters in a string.</p> <p>For example, the regular expression [A-Z] matches the input with the characters A through Z.</p> <p>If the input consists of characters other than A through Z, an error message (as specified in the Invalid Message field) appears.</p>
Invalid Message	Enter an error message that you want displayed on the UI when the input value does not match the specified regular expression.

Event List

Event Name	Select an event from the list based on which the parameter is conditionally displayed.
Event Handler	Enter a JavaScript function that specifies the actions that the event handler takes in response to an event.

What's Next

You can assign the configuration template to device templates or deploy the template on devices; see ["Assign Configuration Templates to a Device Template" on page 292](#) or ["Deploy a Configuration Template to a Device" on page 293](#).

RELATED DOCUMENTATION

[Devices Overview | 130](#)

[Device Templates Overview | 298](#)

[Edit Configuration Templates Assigned to a Device Template | 302](#)

Preview and Render a Configuration Template

You must be an administrator or a user with the preview privilege to preview configuration templates.

You can use the Preview workflow to validate a configuration template by entering values for the configuration template and then render the template to view the configuration.

We recommend that you use this workflow to validate a configuration template before assigning it to a device template or deploying it on a device.

To preview and render a configuration template:

1. Select **Configuration > Templates > Config Templates**.

The Configuration Templates page appears.

2. Select the configuration template that you want to preview and click **Preview**.

The Template Preview for *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 renders 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, Clone, and Delete a Configuration Template](#)" on page 295.

6. Click **Close**.

You are returned to the Configuration Templates page. You can assign the configuration template to one or more device templates or deploy on a device.

RELATED DOCUMENTATION

[Deploy a Configuration Template to a Device | 293](#)

[Configuration Templates Overview | 278](#)

[Device Templates Overview | 298](#)

Assign Configuration Templates to a Device Template

To assign a configuration template to a device template, you must be an administrator or a user with the privilege to assign configuration template to a device template

To assign a configuration template to one or more device templates:

1. Select **Configuration > Templates > Config Templates**.

The Configuration Templates page appears.

2. Select the configuration template that you want to assign and select **More > Assign to Device Template**.

The Assign To Device Templates page appears. This page lists the device templates to which you can assign the configuration template.

3. Click the **Set Initial Configuration** to configure initial values for the parameters defined in the configuration template. These values are configured on the device when you assign the device template to the device during device activation.

The Initial Configuration page appears.

4. Assign initial values for the parameters and click **OK**.

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

You are returned to the Assign to Device Templates page.

5. In the Assign to Device Templates page, under Device Templates, select one or more device templates to which you want to assign the configuration template.
6. Click **OK**.

You are returned to the Configuration Templates page and a popup appears indicating whether the assignment is successful or has failed. If the assignment failed, you can retry the assignment or contact Juniper Networks support.

If the assignment is successful, you can navigate to the Device Templates page (Configuration > Templates > Device Templates) where you can view the configuration templates assigned to a device template; see "[Edit Configuration Templates Assigned to a Device Template](#)" on page 302.

RELATED DOCUMENTATION

[Configuration Templates Overview | 278](#)

[Configuration Templates Workflow | 280](#)

Deploy a Configuration Template to a Device

You can deploy a configuration template directly on one or more devices that were previously activated. This enables you to add configurations to devices after a device was activated or to deploy additional configuration to the device.

To deploy a configuration template on a device, you must either be an administrator or a user with the privilege to deploy configuration on devices.

To deploy a configuration template to one or more devices:

1. Select **Configuration > Templates > Config Templates**.

The Configuration Templates page appears.

2. Select the configuration template that you want to deploy and click **Deploy to Devices**.

The list of devices to which the configuration template can be assigned appears in the Configuration Templates page.

3. Do one of the following:

- If you have not set values for the parameters in the configuration template, click **Set Parameters**.

The Template Parameters page appears.

a. In the Configure tab, assign values for the parameters.

b. Click **Preview** to view and render the configuration.

If the configuration is fine, click **OK** or change the configuration in the Preview tab if you want to change the configuration.

On clicking OK, a message indicating that the configuration is successful appears and you return to the Devices list.

c. (Optional) Click **Validate** to validate the configuration on the device.

A message indicating that a job is created for the validation appears. You can view the status of the validation from the **Monitor > Jobs** page.

- If the values are set for the configuration template, then click **Validate** to validate the configuration on the device.

A message appears indicating that a job is created for the validation. You can view the status of the validation from the **Monitor > Jobs** page.

4. Click **Deploy**.

The Deploy page appears.

5. Do one of the following:

- Click **Deploy Now** to deploy the configuration on the selected devices immediately.
- Click **Deploy Later** to deploy the configuration later.

If you choose to deploy the configuration later, you must enter the date (in MM/DD/YYYY format) and time (in HH:MM:SS 24-hour or AM/PM format) that you want the deployment to occur.

6. Click **OK**.

The settings that you entered are saved and you are returned to the Configuration Templates page. A confirmation message appears indicating that a job is created. For each device, a separate task is triggered in the job to deploy the configuration.

You can view the status of the jobs from the Jobs page (**Monitor > Jobs**).

RELATED DOCUMENTATION

[Configuration Templates Workflow | 280](#)

[Add Configuration Templates | 284](#)

[View and Manage Device Configuration | 165](#)

Edit, Clone, and Delete a Configuration Template

IN THIS SECTION

- [Edit a Configuration Template | 295](#)
- [Clone a Configuration Template | 296](#)
- [Delete a Configuration Template | 296](#)

You should be an administrator or a user with edit, clone, and delete privileges to edit, clone, and delete configuration templates.

Edit a Configuration Template

To edit a Configuration Template:

1. Select **Configuration > Templates > Config 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 same as the fields that you configure in the Add Configuration Template workflow.

3. Modify the fields as needed.

Refer "[Add Configuration Templates](#)" on page 284 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 or assigned to a device template, then you must redeploy the configuration template for the changes to take effect.

Clone a Configuration Template

To clone a configuration template:

1. Select **Configuration > Templates > Config Templates**.

The Configuration Templates page appears.

2. Select the configuration template that you want to clone and click **Clone**.

The Clone Configuration Template page appears.

3. In the **Template Name** field, enter a unique template name that can only contain alphanumeric characters and hyphens up to a maximum of 64 characters.

4. Click **OK**.

You are returned to the Configuration Templates page and a confirmation message appears at the top of the page indicating the status of the clone operation.

After a template is cloned successfully, you can modify the template if needed. See the preceding section for details.

Delete a Configuration Template

NOTE:

- You cannot delete predefined configuration templates.
- You can delete a configuration template only if the following conditions hold good:
 - You added (created) the template.
 - The template is not assigned to a device template.
 - The template is not deployed on a device.

1. Select **Configuration > Templates > Config Templates**.

The Configuration Templates page appears.

2. Select the configuration template that you want to delete and click the **Delete (X)** icon.

You are asked to confirm the delete operation.

3. Click **Yes.**

You are returned to the Configuration Templates page and a popup appears indicating whether the deletion was successful or not.

SEE ALSO

[Configuration Templates Workflow | 280](#)

[Edit, Clone, and Delete a Configuration Template | 295](#)

Device Templates

IN THIS CHAPTER

- [Device Templates Overview | 298](#)
- [About the Device Templates Page | 299](#)
- [Edit Configuration Templates Assigned to a Device Template | 302](#)
- [Edit, Clone, and Delete Device Templates | 303](#)

Device Templates Overview

IN THIS SECTION

- [Predefined Device Templates | 298](#)

Paragon Automation provides device templates that contain configurations for a device. This configuration is applied while onboarding or discovering a device.

Paragon Automation provides several default device templates for Juniper Networks devices. You can either use the default device template provided with Paragon Automation if the template suits your requirements or if you have specific requirements, you can customize the default device template to meet your requirements.

Predefined Device Templates

[Table 63 on page 299](#) lists the predefined device templates provided with Paragon Automation.

Table 63: Predefined Device Templates

Template Name	Description
juniper-ex-default-device-template	Default device template for Juniper Networks EX Series devices.
juniper-mx-default-device-template	Default device template for Juniper Networks MX Series devices.
juniper-qfx-default-device-template	Default device template for Juniper Networks QFX Series devices.

RELATED DOCUMENTATION

[Zero-Touch Provisioning Overview | 132](#)

[Supported Devices | 144](#)

[Configuration Templates Overview | 278](#)

About the Device Templates Page**IN THIS SECTION**

- [Tasks You Can Perform | 299](#)
- [Field Descriptions | 301](#)

To access this page, click **Configuration > Templates > Device Templates**. Use this page to view details of the device templates and manage them.

Tasks You Can Perform

You can perform the following tasks from this page:

- View details of a device template.

To view details of a device template, select the device template and click the **Details** button. The Details for Device Template pane appears on the right side of the Device Templates page, displaying details such as the device family to which the template is applicable, description, number of configuration templates assigned, and date and time the device template was last updated. Click the **Close** icon (X) to close the pane.

- Edit the configuration templates assigned to a device template; see ["Edit Configuration Templates Assigned to a Device Template" on page 302.](#)
- Edit, clone, and delete a device template; see ["Edit, Clone, and Delete Device Templates" on page 303.](#)
- Filter Entries—Filter the table entries by adding new filtering criteria.

Hover over the **Filter** (funnel) icon and select **Add Filter**. On the **Add Criteria** page, select the filtering criteria from the **Field** and **Condition** list, and enter the text to be compared in the **Value** field. Click **Add**.

The filtered table entries are listed and the filter criteria name is displayed above the table column names.

You can click the cross (X) icon (next to the filter name) to remove the filtering criteria.

NOTE:

- You can add multiple filtering criteria. Once you add the multiple filtering criteria, you can select the **And** condition to display the entries matching all the filtering criteria or select the **Or** condition to display the entries matching any one of the filtering criteria.
- Quick filter: Save the filtering criteria as quick filters. Once you have added all the filtering criteria, you can save a particular criteria or multiple criteria for future use by clicking **Save**.

On the **Save Filter** window, 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 a specific column in the table.

Hover over the **More Options** (vertical ellipsis) > **Show/Hide Columns** and select the *Column-Name* check box of the columns you want to display in the table.

- **Reset Preference**—Reset the displayed columns to the default set of columns for each tab in the table.

Hover over the **More Options** (vertical ellipsis) and select **Reset Preference**.

- **Sort Entries**—Click the column name to highlight the up and down arrows next to the column name. Sort the table entries in ascending or descending order of that column by clicking the up or down arrow respectively.

Field Descriptions

Table 64 on page 301 displays the fields on the Device Templates page.

Table 64: Fields on the Device Templates Page

Field	Description
Name	The name of the device template.
Description	A brief description of the device template.
Family	The device family to which the device template can be applied; for example, Juniper-EX, Juniper-MX, and Juniper-QFX
Default	<p>Indicates whether the device template is the default template.</p> <ul style="list-style-type: none"> • True indicates that the device template is the default device template for the device family. • False indicates that the device template is not the default device template for the device family. <p>The default device template for a device family is applied to the devices of that family during device onboarding or discovery.</p>
Last Updated	Date and time when the device template was last updated, in the Month DD, YYYY HH:MM:SS AM/PM format.
Created By	The user who added the device template to Paragon Automation.

RELATED DOCUMENTATION

[Zero-Touch Provisioning Overview | 132](#)

[Supported Devices | 144](#)

[Add Devices Overview | 135](#)

Edit Configuration Templates Assigned to a Device Template

An administrator or a user with permissions to edit device templates only can edit the configuration templates assigned to a device template.

While editing configuration templates assigned to a device template, you can do the following tasks:

- Set initial values for one or more configuration templates. The values that you set initially are assigned to the device during the device discovery or device onboarding.
- Assign one or more configuration templates.
- Remove one or more assigned configuration templates.

To edit a configuration templates assigned to a device template:

1. Select **Configuration > Templates > Device Templates** on the left navigation menu.

The Device Templates page appears.

2. Click on the device template for which you want to edit assigned configuration templates.

The Configuration Templates page appears.

3. Do one of the following:

- To set initial configuration values:

- a. Select a configuration template and click **Set initial Configuration**.

The Initial Configuration page appears.

- b. Set the initial values for the parameters in the configuration template and click **OK**.

Initial values are assigned to the configuration template.

- To add configuration templates to a device template:

- a. Click the **Add (+)** icon.

The Add New Configuration Template page appears listing the available configuration templates in the Templates list.

- b. From the **Templates** list, select the templates to be added, one at a time, and click **OK**.

A confirmation message appears indicating that the configuration template is added to the device template and you are returned to the Configuration Templates page.

- To remove one or more configuration templates assigned to the device template:
 - a. Select the configuration templates that you want to remove and click the **Delete** (trashcan) icon.

A confirmation message appears.

- b. Click **Yes** to confirm.

The selected configuration templates are removed from the device template.

Click **OK**.

The configuration template assignment in the device template is modified and you are returned to the Device Templates page.

RELATED DOCUMENTATION

[Configuration Templates Workflow | 280](#)

[Assign Configuration Templates to a Device Template | 292](#)

[Zero-Touch Provisioning Overview | 132](#)

Edit, Clone, and Delete Device Templates

IN THIS SECTION

- [Edit a Device Template | 304](#)
- [Clone a Device Template | 304](#)
- [Delete Device Templates | 305](#)

To edit, clone, or delete a device template, you must be an administrator or have permissions to edit, clone, and delete device templates.

Edit a Device Template

You can edit only the following parameters of a device template:

- Description
- Enable or disable the template as the default template

To edit a device template:

1. Select **Configuration > Templates > Device Templates** on the left navigation menu.

The Device Templates page appears.

2. Select the device template that you want to edit and click the **Edit** (pencil) icon.

The Edit Device Template page appears.

3. Do the following: and enable or disable the template as the default template.

- Edit the description in the **Description** field.

The description can be a set of alphanumeric characters and some special characters [hyphen (-), underscore (_), and period (.)]. There is no limit on the number of characters used.

- Click the **Mark as Default** toggle button to enable or disable using the device template as the default template for the device family.

NOTE: You can have one default device template per device family.

4. Click **OK**.

A confirmation message appears indicating that the device template is successfully updated and the updates appear on the Device Templates page.

Clone a Device Template

Cloning a device template is useful when you want to create a device template that is similar to an existing one but with small differences.

To clone a device template:

1. Select **Configuration > Templates > Device Templates** on the left navigation menu.

The Device Templates page appears.

2. Select the device template that you want to clone and click the **Clone** button.

The Clone Device Template page appears.

3. Do the following:

- In the **Name** field, edit the name of the device template. The name can be a set of alphanumeric characters, some special characters [hyphen (-), underscore (_), period (.) and space]. There is no limit on the number of characters used.
- In the **Description** field, edit the description of the device template. There is no restriction on the description you provide.
- Click the **Mark as Default** toggle button to enable or disable setting the device template as the default device template.

NOTE: You can have only one default device template per device family. If you set the cloned device template as the default for the device family, the default tag on the earlier device template is removed.

4. Click **OK**.

A message appears indicating that the device template is cloned successfully and listed on the Device Templates page.

Delete Device Templates

To delete a device template:

1. Select **Configuration > Templates > Device Templates**.

The Device Templates page appears.

2. Select the device template that you want to delete and click the **Delete** button.

A confirmation message appears.

3. Click **Yes** to delete the device template.

The device template is deleted and no longer listed on the Device Templates page.

SEE ALSO

[Device Templates Overview | 298](#)

[Zero-Touch Provisioning Overview | 132](#)



Manage Playbook, Rules, Resources, and Graphs

[Playbooks | 307](#)

[Rules | 320](#)

[Resources | 369](#)

[Graphs | 401](#)

[Grafana | 423](#)

Playbooks

IN THIS CHAPTER

- [About Playbooks | 307](#)
- [Add a Predefined Playbook | 308](#)
- [Create a Playbook Using the Paragon Insights GUI | 309](#)
- [Edit a Playbook | 310](#)
- [Clone a Playbook | 311](#)
- [Manage Playbook Instances | 312](#)

About Playbooks

In order to fully understand any given problem or situation on a network, it is often necessary to look at a number of different system components, topics, or key performance indicators (KPIs). Paragon Insights operates on playbooks, which are collections of rules for addressing a specific use case. **Playbooks** are the Paragon Insights element that gets applied, or run, on your device groups or network groups.

Paragon Insights comes with a set of pre-defined **Playbooks**. For example, the system-KPI playbook monitors the health of system parameters such as system-cpu-load-average, storage, system-memory, process-memory, etc. It then notifies the operator or takes corrective action in case any of the KPIs cross pre-set thresholds. Following is a sample list of Juniper-supplied Playbooks.

- bgp-session-stats
- route-summary-playbook
- lldp-playbook
- interface-kpis-playbook
- system-kpis-playbook
- linecard-kpis-playbook
- chassis-kpis-playbook

You can create a playbook and include any rules you want in it. You apply these playbooks to device groups. By default, all rules contained in a Playbook are applied to all of the devices in the device group. There is currently no way to change this behavior.

If your playbook definition includes network rules, then the playbook becomes a network playbook and can only be applied to network groups. You must not add device group rules and network rules in a single playbook.

To access the Playbooks page in the Paragon Automation Platform GUI, go to **Configuration > Playbooks**.

You can perform the following tasks in Playbooks page.

- Add a Pre-Defined Playbook. See "[Add a Predefined Playbook](#)" on page 308.
- Create a New Playbook Using the Paragon Automation GUI. See "[Create a New Playbook Using the Paragon Insights GUI](#)" on page 309.
- Edit a Playbook. See "[Edit a Playbook](#)" on page 310.
- Clone a Playbook. See "[Clone a Playbook](#)" on page 311.
- Manage Playbook Instances. See "[Manage Playbook Instances](#)" on page 312.

RELATED DOCUMENTATION

| [Rules Overview](#) | 321

Add a Predefined Playbook

To add a predefined playbook to Paragon Insights:

1. In a browser, go to <https://github.com/Juniper/healthbot-rules> and download the predefined playbook file to your computer.
2. In the Paragon Insights GUI, click the **Configuration > Playbooks** icon in the Paragon Automation menu.
3. Click the upload playbook button (↑ **Upload Playbook**).
4. Click the **Choose Files** button.
5. Navigate to the playbook file and click **Open**.
The extension of playbook files must be .playbook.
6. Select one of the following options:

Upload	Upload the file and add the playbook but do not deploy the updated configuration. You can use this option when, for example, you are making several changes and want to deploy all your updates at the same time.
Upload & Deploy	Upload the file, add the playbook, and immediately deploy the configuration.

RELATED DOCUMENTATION

[Manage Playbook Instances | 312](#)

[Create a Playbook Using the Paragon Insights GUI | 309](#)

Create a Playbook Using the Paragon Insights GUI

Paragon Insights operates on playbooks, which are a collection of rules for solving a specific customer use case. For example, the system-kpi-playbook monitors the health of system parameters such as the system-cpu-load-average, storage, system-memory, process-memory, and so on. The playbook also notifies the operator or takes corrective action in case any of the KPIs cross pre-set thresholds. Any single rule can be a part of 0, 1, or more playbooks. The playbook is the element that gets deployed on device groups. Rules that you do not include in any playbook are not deployed to any device.

NOTE: Click the **Name**, **Running**, **Paused**, or **Synopsis** column headers in the Playbooks table to organize the data in ascending or descending order.

To create a playbook using the Paragon Insights GUI:

1. Click the **Configuration>Playbooks** icon in the left navigation bar.
2. Click the create playbook button (+ **Create Playbook**).
A new window appears with 4 fields: Name, Synopsis, Description, and Rules. We describe the use of each field in the following steps.
3. Enter a name for the playbook in the **Name** field.
4. Enter a short description for the playbook in the **Synopsis** field.
This text appears in the Synopsis column of the table on the **Playbooks** page.
5. (Optional) Enter a description of each rule added in this playbook in the **Description** field.
This text can only be seen if you click on the playbook name on the **Playbooks** page.

6. From the rules list, select the rules that make up this playbook.
7. Select one of the following options:

Save	Save and add the playbook but do not deploy the updated configuration. You can use this option when, for example, you are making several changes and want to deploy all your updates at the same time.
Save & Deploy	Immediately deploy the configuration and save and add the playbook.

RELATED DOCUMENTATION

[Add a Predefined Playbook | 308](#)

[Clone a Playbook | 311](#)

Edit a Playbook

To edit a playbook:

1. Click the **Configuration > Playbooks** icon in the left-nav bar.
2. Click the name of the playbook.
3. Modify the necessary text boxes.
4. Click **Save** to only save the configuration.

Click **Save and Deploy** to save and deploy the configuration immediately.

(Optional) You can delete a playbook by clicking the trash can icon in the **Delete** column.

NOTE: You cannot edit or delete a system defined (Juniper provided) playbook.

When you update a playbook, the new changes in the playbook are not applied to the existing instances of the playbook. For example, a playbook instance that is associated to a device group will not be updated when the playbook is edited or updated. You must delete the existing playbook instance and create a new one for updates to be applied.

RELATED DOCUMENTATION

[Manage Playbook Instances](#) | 312

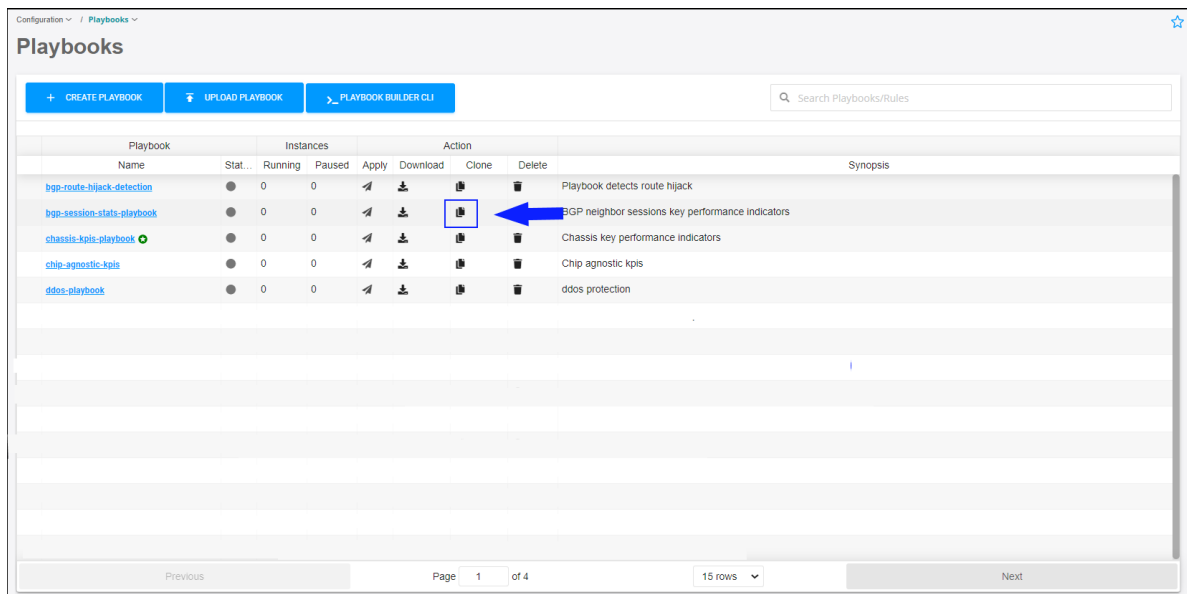
Clone a Playbook

You can clone an existing playbook and modify configurations.

Follow these steps to clone a playbook by using the Paragon Insights UI.

1. Click **Configuration**>**Playbooks** in the left-nav bar.
The Playbooks page is displayed.
2. Click the **Clone** icon as show in [Figure 15 on page 311](#).

Figure 15: Clone a Playbook



The Clone Playbook: *<name of playbook>* page is displayed.

You can now edit **Name**, **Synopsis**, and **Description** sections. You can also add new rules or remove existing rules from the Rules drop-down list.

3. Click **Save** to save configurations and confirm cloning the playbook.
Alternatively, click **Save & Deploy** to save configurations, confirm cloning the playbook, and deploy the newly cloned playbook.

RELATED DOCUMENTATION

[Edit a Playbook | 310](#)

Manage Playbook Instances

IN THIS SECTION

- [View Information About Playbook Instances | 312](#)
- [Create and Run a Playbook Instance | 314](#)
- [Manually Pause or Play a Playbook Instance | 316](#)
- [Create a Schedule to Automatically Play/Pause a Playbook Instance | 317](#)

View Information About Playbook Instances

To view information about playbook instances:

1. Click the **Configuration > Playbooks** option in the Paragon Automation menu.

You can see the saved playbooks on the main Playbooks page.

Playbook	Instances		Action			Synopsis
	Name	Running	Paused	Apply	Live	
bgp-route-hijack-detection	0	0	↗	●	🗑️	Playbook detects route hijack
bgp-session-stats-playbook 🟢	0	0	↗	●	🗑️	BGP neighbor sessions key performance indicators
chassis-kpis-playbook 🟢	0	0	↗	●	🗑️	Chassis key performance indicators
chip-agnostic-kpis	0	0	↗	●	🗑️	Chip agnostic kpis
DHCP-server-statistics 🟢	0	0	↗	●	🗑️	DHCP Local Server and Relay Statistics KPIs
dot1x-user-authentication-kpis	0	0	↗	●	🗑️	dot1x authentication KPI
evpn-irb-icmp-probe	0	0	↗	●	🗑️	EVPN-VXLAN n/i key performance indicators
evpn-vxlan-kpis 🟢	0	0	↗	●	🗑️	EVPN-VXLAN key performance indicators
forwarding-table-summary 🟢	0	0	↗	●	🗑️	Forwarding table and protocol routes key performance indicators
get-vpn-stats Requires advanced licensing	0	0	↗	●	🗑️	Interface and routing instance collector
icmp-outlier 🟢	0	0	↗	●	🗑️	ICMP outlier detector
icmp-probe 🟢	0	0	↗	●	🗑️	ICMP RTT response checker
▶ interface-kpis-playbook	1	0	↗	●	🗑️	Interface key performance indicators
interface-optical-kpis	0	0	↗	●	🗑️	Optical interface key performance indicators
isis-stats-playbook	0	0	↗	●	🗑️	ISIS adjacency key performance indicators

Page 1 of 3 | 15 rows

- Playbooks, applied to a device group or a network group, have a right arrow next to the playbook name.

- The Instances column in the table shows the number of playbook instances running and paused.
- Some playbooks require the purchase and installation of advanced or premium licenses. These playbooks have a green circle with a white star in it. As shown above, it tells you which license you need when you hover your mouse over the icon.
- The Live column (in the Action section of the table) shows a colored circle indicator that represents the overall status of the playbook instances for each playbook. The following table provides the color definitions:

Table 65: Color Definitions for the Live Column

Color	Definition
Green	All instances associated with this playbook are currently running.
Yellow	One or more instances associated with this playbook are paused.
Gray/Black	Either, no instance is available for this playbook or an instance is available for this playbook but the configuration is not deployed yet.

2. Click on the right arrow next to the playbook name to expand or collapse the playbook instance details. If no caret is present, then the playbook is not applied to any device groups or network groups.

Playbook instances have the following details:

Column Name or Widget	Description
Instance Name	User-defined instance name.
Schedule	Name of the schedule profile applied to the playbook instance. For information on how to configure a schedule profile, see "Create a Schedule to Automatically Play/Pause a Playbook Instance" on page 317 Click on the name to display the schedule details.
Device/Network Group	The Device group or the network group on which a scheduler is running.

(Continued)

Column Name or Widget	Description
No. of devices	Number of devices that run this playbook instance. This column is applicable for device group instances only, not for network group instances.
Status	<p>Current status of the playbook instance. Status can be either Running or Paused. The Status column also indicates whether the action was performed automatically or manually.</p> <p>Note: If the status of a playbook instance is Running (manual), you can manually pause the instance using the Pause Instance button. In this case, the status will change to Paused (manual). To resume running the schedule for this instance, you must manually run the instance using the Play Instance button. In this case, the status will change back to Running (manual).</p>
Started/Paused at	Date and time when the playbook instance was last started or paused. The date reflects local browser time zone.
Next Action	This column applies only to playbook instances associated with a schedule. It indicates whether the playbook instance is scheduled to automatically pause or play in the future. This column is blank if no schedule runs on the playbook instance or if the status of the instance is Paused (manual) .
Play/Pause button	<p>Pauses or plays a playbook instance or the schedule for a playbook instance.</p> <p>The Play/Pause button toggles between the two states. For more information, see "Manually Pause or Play a Playbook Instance" on page 316.</p>
Trash can icon	Deletes the playbook instance.

Create and Run a Playbook Instance

To create a playbook instance for a device group:

1. Click the **Configuration > Playbooks** option in the left navigation bar.
2. Click the **Apply** icon (in the Action section of the table) for the desired playbook.

A pane titled *Run Playbook: <playbook-name>* appears.

3. In the **Name of Playbook Instance** field, you must enter an appropriate name for this instance of the playbook.
4. (Optional) In the **Run on schedule** field, choose the name of the schedule that you want to apply to this playbook instance. You can apply only one schedule for a playbook instance. If you want a specific playbook to run multiple schedules, you must create multiple instances. Each instance must have its own unique name and schedule.

For information on how to configure a schedule, see ["Create a Schedule to Automatically Play/Pause a Playbook Instance" on page 317](#).

To view information about an existing scheduler:

- a. Click the **Configuration > Insights Settings** in the Paragon Automation menu.
 - b. In the **Scheduler Settings**, you can see a summary of the properties for each saved schedule in the table. Click on a specific schedule name to view additional details.
5. In the **Device Group** section under **Rules**, apply this playbook instance to the appropriate device group using the list.

The list of devices in the **Devices** section changes based on the device group selected.

NOTE: If your playbook contains network rules, the **Device Group** section does not appear. Instead, you configure the **Network Groups** section (not shown).

6. Click one of the devices listed in the **Devices** section.
Here is where you can customize the variables (defined in rules added in the playbook) for this device.
7. In the section titled **Variable values for Device <Device Name>**, you can see various parameters for each rule associated with the playbook. The default values for each variable are displayed as Gray text in the fields. You can leave these values unchanged or override them by entering a new value.
Repeat steps 6 and 7 for each device in the device group, as needed.

When you create an instance of the default **icmp-outlier** playbook, you earlier had to enter the round trip time (RTT) XPath for each device (using device id). Paragon Automation supports splat operator in ICMP outlier detection playbook. You can enter the splat operator (*) in regular expression format in the playbook so that, the rules are applied to all devices in the device group.

For example:

```
/device-group[device-group-name=core]/device[device-id=~/.*/]/topic[topic-name=protocol.icmp]/rule[rule-name=check-icmp-statistics]/rtt-average-ms
```

You can also use the splat operator when you add a playbook instance with a custom network rule for outlier detection.

NOTE: If you delete or add a device in a device group on which an ICMP outlier playbook instance runs, you must pause the playbook instance, modify the XPath configuration to reflect the addition or deletion of devices, and re-apply the playbook instance for the device group.

8. When you are satisfied that all of the variable values are appropriate for all the devices in the device group, select one of the following options.

Save	Save your edits but do not deploy the updated configuration and do not run the instance. You can use this option when, for example, you are making several changes and want to deploy all your updates at the same time.
Save & Deploy	Deploy the configuration and immediately run the instance. If you apply a schedule profile, the instance runs according to the configuration of the profile.

Manually Pause or Play a Playbook Instance

When you pause an instance, Paragon Automation does not collect, analyze or raise an alarm for the device or the network groups associated with the playbook rules. Data collected before pausing the instance remains in the system, but Paragon Automation does not collect or analyze new data until you play the instance again.

The following table describes the state of the playbook instance in the Play/Pause column:

If the displayed Play/Pause button is...	Then the state of the playbook instance is...
Pause Instance	<ul style="list-style-type: none"> • The instance is running. • The instance is not associated with a schedule. • Paragon Automation is collecting the data.
Play Instance	<ul style="list-style-type: none"> • The instance is not running. • The instance is not associated with a schedule. • Paragon Automation is not collecting the data.

To manually pause a playbook instance:

1. Click the **Configuration > Playbooks** option in the left navigation bar.
2. Click the right arrow next to the name of the playbook that you want to pause.
3. Click the **Pause Instance** button to pause a playbook instance (not associated with a schedule).
4. The Play/Pause Playbook Instance dialog box appears. Select one of the following options:

Pause	Flags this playbook instance to be paused the next time you deploy the configuration. Use this option if you are making several changes and want to deploy all your edits at the same time.
Pause & Deploy	<p>Immediately pause the playbook instance and deploy the configuration. It will take a few seconds for the playbook table to update to show that the instance is paused.</p> <p>You see a slight delay in status updates in the table because the play and the pause actions are asynchronous. You can track the status of this asynchronous activity through the deploy icon located in the upper right corner of the window (as indicated in the success message of deploy action). Once this action is complete, the status is reflected in the playbook table as well.</p>

Once the application refreshes the playbook table, the playbook name shows a yellow icon in the Live column as a visual indicator that an instance is paused.

5. To resume a paused playbook instance, click the **Play Instance** button to resume running a playbook instance (not associated with a schedule).

Create a Schedule to Automatically Play/Pause a Playbook Instance

To automatically play/pause a playbook instance, you must first create a scheduler and then apply the scheduler to the playbook instance. You can apply only one schedule for a playbook instance. If you want a specific playbook instance to run on multiple schedules, you must create multiple versions of the instance, each with its own unique name and scheduler.

To create a schedule for a playbook instance:

1. Click the **Configuration > Insights Settings** option in the left navigation bar.
2. Click the **Scheduler** tab.
3. In **Scheduler Settings**, click the add scheduler button (**+ Scheduler**).
4. Enter the necessary values in the text boxes and select the appropriate options for the playbook instance schedule.

The following table describes the attributes in the **Add a scheduler** and **Edit a scheduler** panes:

Attributes	Description
Name	Enter the name of the playbook instance scheduler.

(Continued)

Attributes	Description
Scheduler Type	<p>Choose discrete or continuous.</p> <p>For discrete schedulers, you can configure a discrete length of time to play the playbook instance using the Run for field. Once the run time has ended, Paragon Insights will automatically pause the instance. You can also configure Paragon Insights to automatically resume playing the instance using the Repeat field.</p> <p>For continuous schedulers, you can configure Start on, End on, and Repeat fields.</p>
Start On	<p>Use the pop-up calendar to select the date and time to play the playbook instance for the first time.</p>
Run for	<p>Configure a time period for the discrete playbook instance. First enter an integer value and then choose the unit of measure (minute, hour, or day) from the list.</p> <p>Once the run time has ended, Paragon Insights will automatically pause the instance. You can also configure Paragon Insights to automatically resume playing the instance using the Repeat field.</p>
End On	<p>(Optional) Use the pop-up calendar to select the date and time to pause the playbook instance. Leave blank if you want the playbook instance to play indefinitely.</p> <p>For discrete schedulers, you can configure Run for time and End on time. If the End on time period is shorter than Run for time, the discrete scheduler stops running after the End on time.</p>
Repeat	<p>Configure the Run for field before configuring the Repeat field. The Custom Repeat interval you enter must be larger than the configured Run for length of time.</p> <p>In the list, choose one of the following:</p> <ul style="list-style-type: none"> • The frequency (day, week, month, or year) at which you want the playbook instance to play. • The Never option if you want the playbook to play only once. • The Custom option to specify a custom frequency at which you want the playbook instance to play. Use the Repeat Every field to configure the custom frequency.

(Continued)

Attributes	Description
Repeat Every	(Optional) If you chose the Custom option for the Repeat field, enter the custom frequency at which you want the playbook instance to play. First enter an integer value and then choose the unit of measure (minute, hour, or day) from the list.

5. Do one of the following:
 - **Save** – Click **Save** to save the scheduler configuration but not deploy it in Paragon Automation.
 - **Save and Deploy** – Click **Save and Deploy** to save and deploy the configuration in your Paragon Automation application.
6. Now you're ready to apply the scheduler to a playbook instance. See "[Create and Run a Playbook Instance](#)" on page 314 for more information.

RELATED DOCUMENTATION

| [About the Rules Page](#) | 340

Rules

IN THIS CHAPTER

- Understand Paragon Insights Topics | 320
- Rules Overview | 321
- About the Rules Page | 340
- Add a Predefined Rule | 340
- Edit, Clone, Delete, and Download Rules | 341
- Configure a Custom Rule in Paragon Automation GUI | 343
- Configure Paragon Insights Notification for LSP Gray Failures | 359
- Configure Multiple Sensors per Device | 362
- Understand Sensor Precedence | 364
- Configure Sensor Precedence | 365

Understand Paragon Insights Topics

Network devices are made up of a number of components and systems from CPUs and memory to interfaces and protocol stacks and more. In Paragon Insights, topics are the construct used to address those different device components. The **Topics** block is used to create name spaces that define what needs to be modeled. The **Topics** block consists of **Rules** blocks which in turn consist of the **Fields** blocks, **Functions** blocks, **Triggers** blocks, etc. See "[Rules Overview](#)" on page 321. Each rule created in Paragon Automation Platform must be included in a *Topic*. Juniper has curated a number of these system components into a list of *Topics*:

- chassis
- class-of-service
- external
- firewall
- interfaces

- kernel
- linecard
- logical-systems
- protocol
- routing-options
- security
- service
- system

Any pre-defined rules provided by Juniper fit within one of these topics with the exception of *external*. The *external* topic is reserved for user-created rules. You can create sub-topics underneath any of the allowed topic names by appending *.<sub-topic>* to them. For example, *kernel.tcpip*,

In the Paragon Automation GUI, when you create a new rule, the *Topic* field is automatically populated with the *external* topic name.

RELATED DOCUMENTATION

[Add a Predefined Rule | 340](#)

[Configure a Custom Rule in Paragon Automation GUI | 343](#)

Rules Overview

IN THIS SECTION

- [Rules | 324](#)
- [Sensors | 329](#)
- [Fields | 330](#)
- [Vectors | 332](#)
- [Variables | 333](#)
- [Functions | 334](#)
- [Triggers | 335](#)


- Tagging | 338
- Rule Properties | 338
- Pre/Post Actions | 338

A rule is a package of components, or blocks, needed to extract specific information from the network or from a Junos device. **Rules** conform to a specifically tailored domain specific language (DSL) for analytics applications. The DSL is designed to allow rules to capture:

- The minimum set of input data that the rule needs to be able to operate
- The fields of interest from the configured sensors
- The reporting or polling frequency
- The set of triggers that operate on the collected data
- The conditions or evaluations needed for triggers to kick in
- The actions or notifications that need to be performed when a trigger kicks in

The structure of a rule looks like this:

Figure 16: Rule Structure



To keep rules organized, Paragon Insights organizes them into *topics*. Topics can be very general, like **system**, or they can be more granular, like **protocol.bgp**. Each topic contains one or more rules. To know more about topics, see "[Understand Paragon Insights Topics](#)" on page 320.

The details around rules are presented in the following sections.

Rules

Rules are meant to be free of any hard coding. Think of threshold values; If a threshold is hard coded, there is no easy way to customize it for a different customer or device that has different requirements. Therefore, rules are defined using parameterization to set the default values. This allows the parameters to be left at default or be customized by the operator at the time of deployment. Customization can be done at the device group or individual device level while applying the Playbook in which the individual rules are contained.

Rules that are device-centric are called device rules. Device components such as chassis, system, line cards, and interfaces are all addressed as Topics in the rule definition. Generally, device rules make use of sensors on the devices.

Rules that span multiple devices are called network rules. Network rules:

- Must have a rule-frequency configured
- Must not contain sensors
- Cannot be mixed with device rules in a playbook

To deploy either type of rule, include the rule in a playbook and then apply the playbook to a device group or network group.

NOTE: Paragon Insights comes with a set of pre-defined rules.

Not all of the blocks that make up a rule are required for every rule. Whether or not a specific block is required in a rule definition depends on what sort of information you are trying to get to. Additionally, some rule components are not valid for network rules. [Table 66 on page 325](#) lists the components of a rule and provides a brief description of each one.

Table 66: Rule Components

Block	What it Does	Examples	Required in Device Rules?	Valid for Network Rules?
"Sensors" on page 329	<p>The <i>sensor</i> defines the parameters for collecting the data. This typically includes which data collection method to use, which data to ingest, and how often to push or pull the data. In any given rule, a sensor can be referenced by the Fields defined directly within the rule or, it can be referenced from another rule.</p> <p>There are multiple types of sensors available in Paragon Insights : OpenConfig, Native GPB, iAgent, SNMP, and syslog.</p> <p>OpenConfig and iAgent sensors require that a frequency be set for push interval or polling interval respectively. SNMP sensors also require you to set a frequency.</p>	<p>Using the SNMP sensor, poll the network device every 60 seconds to collect all the device data in the Juniper SNMP MIB table jnxOperatingTable.</p>	<p>No-Rules can be created that only use a field reference from another rule or a vector with references from another rule. In these cases, rule-frequency must be explicitly defined.</p>	<p>No</p>

Table 66: Rule Components *(Continued)*

Block	What it Does	Examples	Required in Device Rules?	Valid for Network Rules?
"Fields" on page 330	<p>The <i>fields</i> provide a way to filter or manipulate sensor data, allowing you to identify and isolate the specific pieces of information. Fields can also act as placeholder values, like a static threshold value, to help the system perform data analysis.</p> <p>The source for the Fields block can be a pointer to a sensor, a reference to a field defined in another rule, a constant, or a formula. The field can be a string, integer or floating point. The default field type is string.</p>	Extract, isolate, and store the jnxOperating15MinLoadAvg (CPU 15-minute average utilization) value from the SNMP table specified above in the sensor.	Yes- Fields contain the data on which the triggers operate. Regular fields and key-fields can be added to rules based on conditional tagging profiles. See the "Tagging" on page 338 section below.	Yes
"Vectors" on page 332	The Vectors block allows handling of lists, creating sets, and comparing elements amongst different sets. A vector is used to hold multiple values from one or more fields. <i>Vectors</i> allow you to leverage existing elements to avoid the need to repeatedly configure the same elements across multiple rules.	A rule with a configured sensor, plus a vector to a second sensor from another rule; a rule with no sensors, and vectors to fields from other rules.	No	Yes

Table 66: Rule Components (Continued)

Block	What it Does	Examples	Required in Device Rules?	Valid for Network Rules?
"Variables" on page 333	<p>The Variables block allows you to pass values into rules. Invariant rule definitions are achieved through mustache-style templating like <code>{{<placeholder-variable>}}</code>. The placeholder-variable value is set in the rule by default or can be user-defined at deployment time.</p>	<p>The string "ge-0/0/0", used within a field collecting status for all interfaces, to filter the data down to just the one interface; an integer, such as "80", referenced in a field to use as a static threshold value.</p>	No	No
"Functions" on page 334	<p>The Functions block allows you to extend fields, triggers, and actions by creating prototype methods in external files written in languages like python. The functions block includes details on the file path, method to be accessed, and any arguments, including argument description and whether it is mandatory.</p>	<p>A rule that monitors input and output packet counts, using a function to compare the count values; a rule that monitors system storage, invoking a function to cleanup temp and log files if storage utilization goes above a defined threshold</p>	No	No

Table 66: Rule Components (Continued)

Block	What it Does	Examples	Required in Device Rules?	Valid for Network Rules?
"Triggers" on page 335	<p><i>Triggers</i> periodically bring together the fields with other elements to compare data and determine current device status.</p> <p>The Triggers block operates on fields and are defined by one or more <i>Terms</i>. A trigger Term includes one or more 'when-then' statements, which contain the parameters that define how device status is visualized on the health pages. When the conditions of a Term are met, then the action defined in the Term is taken.</p> <p>By default, triggers are evaluated every 10 seconds, unless explicitly configured for a different frequency.</p> <p>By default, all triggers defined in a rule are evaluated in parallel.</p>	<p>Every 90 seconds, check the CPU 15min average utilization value, and if it goes above a defined threshold, set the device's status to red on the device health page and display a message showing the current value.</p>	<p>Optional—Triggers enable rules to take action.</p>	<p>Yes</p>
"Rule Properties" on page 338	<p>The Rule Properties block allows you to specify metadata for a Paragon Insights rule, such as hardware dependencies, software dependencies, and version history.</p>	<p>Configuration to set the minimum supported release version for backward compatibility for all devices in the MX series of routers.</p>	<p>No</p>	<p>Yes</p>

Table 66: Rule Components (Continued)

Block	What it Does	Examples	Required in Device Rules?	Valid for Network Rules?
"Pre/Post Actions" on page 338	The Pre/Post Actions block allows you to use Action Engine Workflow tasks in a rule such that the Paragon application executes the tasks at the same time as the playbook that contains the rule, or after you stop the playbook instance.	Select a pre-configured action engine workflow as a Pre-Action task that can configure set commands and commit those commands on devices to enable them to use an ingest.	No	No

Sensors

When defining a sensor, you must specify information such as sensor name, sensor type and data collection frequency. As mentioned in [Table 66 on page 325](#), sensors can be one of the following:

- **OpenConfig** For information on OpenConfig JTI sensors, see the [Junos Telemetry Interface User Guide](#).
- **Native GPB** For information on Native GPB JTI sensors, see the [Junos Telemetry Interface User Guide](#).
- **iAgent** The iAgent sensors use NETCONF and YAML-based PyEZ tables and views to fetch the necessary data. Both structured (XML) and unstructured (VTY commands and CLI output) data are supported. For information on Junos PyEZ, see the [Junos PyEZ Documentation](#).
- **SNMP** Simple Network Management Protocol.
- **syslog** system log
- **BYOI** Bring your own ingest – Allows you to define your own ingest types.
- **Flow** NetFlow traffic flow analysis protocol

- **sFlow** sFlow packet sampling protocol

When different rules have the same sensor defined, only one subscription is made per sensor. A key, consisting of *sensor-path* for OpenConfig and Native GPB sensors, and the tuple of *file* and *table* for iAgent sensors is used to identify the associated rule.

When multiple sensors with the same *sensor-path* key have different frequencies defined, the lowest frequency is chosen for the sensor subscription.

Fields

There are four types of field sources, as listed in [Table 66 on page 325](#). [Table 67 on page 330](#) describes the four field ingest types in more detail.

Table 67: Field Ingest Type Details

Field Type	Details
Sensor	<p>Subscribing to a sensor typically provides access to multiple columns of data. For instance, subscribing to the OpenConfig interface sensor provides access to a bunch of information including counter related information such as:</p> <p>/interfaces/counters/tx-bytes, /interfaces/counters/rx-bytes, /interfaces/counters/tx-packets, /interfaces/counters/rx-packets, /interfaces/counters/oper-state, etc.</p> <p>Given the rather long names of paths in OpenConfig sensors, the Sensor definition within Fields allows for aliasing, and filtering. For single-sensor rules, the required set of Sensors for the Fields table are programmatically auto-imported from the raw table based on the triggers defined in the rule.</p>

Table 67: Field Ingest Type Details (Continued)

Field Type	Details
Reference	<p>Triggers can only operate on Fields defined within that rule. In some cases, a Field might need to reference another Field or Trigger output defined in another Rule. This is achieved by referencing the other field or trigger and applying additional filters. The referenced field or trigger is treated as a stream notification to the referencing field. References aren't supported within the same rule.</p> <p>References can also take a time-range option which picks the value, if available, from the time-range provided. Field references must always be unambiguous, so proper attention must be given to filtering the result to get just one value. If a reference receives multiple data points, or values, only the latest one is used. For example, if you are referencing a the values contained in a field over the last 3 minutes, you might end up with 6 values in that field over that time-range. Paragon Insights only uses the latest value in a situation like this.</p> <p>The syntax for reference is:</p> <pre data-bbox="444 892 1058 1024">/device-group[device-group-name=<device-group>]\ /device[device-name=<device>]/topic[topic-name=<topic>]\ /rule[rule-name=<rule>]/field[<field-name>=<field-value>\ AND OR ...]/<field-name></pre> <p>NOTE: The device-group and device components in above path are applicable only for network rules.</p> <p>For example, the pre-defined rule protocol.l3vpn/check-l3vpn-ospf-state has a field to check the interface status using the following reference path:</p> <pre data-bbox="444 1323 1006 1560">/device-group[device-group-name={{pe-device-group}}] /device[device-id={{pe-device-name}}]/ topic[topic-name='protocol.l3vpn'] /rule[rule-name=get-interface-details] /field[sub-interface-index='{{pe-if1-number}}' and interface-name='{{pe-interface-name}}'] /link-state</pre>
Constant	<p>A field defined as a constant is a fixed value which cannot be altered during the course of execution. Paragon Insights Constant types can be strings, integers, and doubles.</p>

Table 67: Field Ingest Type Details (Continued)

Field Type	Details
Formula	<p>Raw sensor fields are the starting point for defining triggers. However, Triggers often work on derived fields defined through formulas by applying mathematical transformations.</p> <p>Formulas can be pre-defined or user-defined (UDF). Pre-defined formulas include: Min, Max, Mean, Sum, Count, Rate of Change, Elapsed Time, Standard Deviation, Microburst, Dynamic Threshold, Anomaly Detection, Outlier Detection, and Predict.</p> <p><para>Rate of Change refers to the difference between current and previous values over their points of time. Packet transfer is an example use case where the Rate of Change formula can be used.</p> <p>The Hold Time field takes a threshold of time interval. The time interval between current and previous values cannot exceed the specified Hold Time value. The Multiplication Factor field is used to convert the unit of the field value. If the field value is calculated in Bytes, specifying 1024 as Multiplication Factor would convert the result into Kilobytes. Hold Time and Multiplication Factor are not mandatory fields when you apply Rate of Change formula.</p> <p>In Paragon Automation, you can get the current point time in Elapsed Time formula by using \$time.</p> <p>Some pre-defined formulas can operate on time ranges in order to work with historical data. If a time range is not specified, then the formula works on current data, specified as <i>now</i>.</p>

Vectors

Vectors are useful in helping to gather multiple elements into a single rule. For example, using a vector you could gather all of the interface error fields. The syntax for Vector is:

```
vector <vector-name>{
  path [$field-1 $field-2 .. $field-n];
  filter <list of specific element(s) to filter out from vector>;
  append <list of specific element(s) to be added to vector>;
}
```

\$field-n can be field of type reference.

The fields used in defining vectors can be direct references to fields defined in other rules:

```
vector <vector-name>{
    path [/device-group[device-group-name=<device-group>]\
/device[device-name=<device>]/topic[topic-name=<topic>]\
/rule[rule-name=<rule>]/field[<field-name>=<field-value>\
AND|OR ...]/<field-name> ...];
    filter <list of specific element(s) to filter out from vector>;
    append <list of specific element(s) to be added to vector>;
}
```

This syntax allows for optional filtering through the <field-name>=<field-value> portion of the construct. Vectors can also take a time-range option that picks the values from the time-range provided. When multiple values are returned over the given time-range, they are all selected as an array.

The following pre-defined formulas are supported on vectors:

- unique @vector1—Returns the unique set of elements from vector1
- @vector1 and @vector2—Returns the intersection of unique elements in vector1 and vector2.
- @vector1 or @vector2—Returns the total set of unique elements in the two vectors.
- @vector1 unless @vector2—Returns the unique set of elements in vector-1, but not in vector-2

Variables

Variables are defined during rule creation on the **Variables** page. This part of variable definition creates the default value that gets used if no specific value is set in the device group or on the device during deployment. For example, the check-interface-status rule has one variable called interface_name. The value set on the **Variables** page is a regular expression (regex), .*, that means all interfaces.

If applied as-is, the check-interface-status rule would provide interface status information about all the interfaces on all of the devices in the device group. While applying a playbook that contains this rule, you could override the default value at the device group or device level. This allows you flexibility when applying rules. The order of precedence is device value overrides device group value and device group value overrides the default value set in the rule.

BEST PRACTICE: It is highly recommended to supply default values for variables defined in device rules. All Juniper-supplied rules follow this recommendation. Default values must not be set for variables defined in network rules.

Functions

Functions are defined during rule creation on the **Functions** tab. Defining a function here allows it to be used in **Formulas** associated with **Fields** and in the **When** and **Then** sections of **Triggers**. Functions used in the when clause of a trigger are known as user-defined functions. These must return true or false. Functions used in the then clause of a trigger are known as user-defined actions.

In Paragon Automation, you can use a Python user-defined function to return multiple values.

For example, consider that you have a function **example_function.py** that has three return values. When you call the **example_function.py** in a rule, the first return value in the user-defined function (UDF) is stored in the rule field (Fields tab) that calls the function. You only need to configure return fields, such as *r2* and *r3*, for the remaining two return values in the Return List of the Functions tab.

In the time-series database, the name of Return List fields are prefixed with the name of the rule field that uses the UDF. For example, *rule_field_name-r2*.

[Figure 17 on page 334](#) shows an example configuration in the **Functions** block. The field details are discussed below.

Figure 17: The Functions Block

The screenshot shows the 'Functions' tab in the Paragon Automation interface. The function being configured is named 'used_percentage'. The configuration includes the following fields and options:

- Function name ***: used_percentage
- Path to Function ***: used-percentage.py
- Method Name ***: used_percentage
- Description**: Enter a description for this function
- Arguments**:
 - Name**: total, **Mandatory**:
 - Name**: used, **Mandatory**:
- Return List**:
 - Name**: used_micro, **Type**: Integer

Buttons for '+ ADD FUNCTION', '+ ADD ARGUMENT', and '+ ADD RETURN LIST' are visible at the bottom of the configuration area.

Triggers

Triggers play a pivotal role in Paragon Insights rule definitions. They are the part of the rule that determines if and when any action is taken based on changes in available sensor data. Triggers are constructed in a when-this, then-that manner. As mentioned earlier, trigger actions are based on **Terms**. A **Term** is built with *when* clauses that watch for updates in field values and *then* clauses that initiate some action based on what changed. Multiple **Terms** can be created within a single trigger.

Evaluation of the *when* clauses in the **Terms** starts at the top of the list of terms and proceeds to the bottom. If a *term* is evaluated and no match is made, then the next *term* is evaluated. By default, evaluation proceeds in this manner until either a match is made or the bottom of the list is reached without a match.

Pre-defined operators that can be used in the *when* clause include:

NOTE: For evaluated equations, the left-hand side and right-hand side of the equation are shortened to LHS and RHS, respectively in this document.

- *greater-than*—Used for checking if one value is greater than another.
 - Returns: True or False
 - Syntax: greater-than <LHS> <RHS> [time-range <range>]
 - Example: //Memory > 3000 MB in the last 5 minutes
 when greater-than \$memory 3000 time-range 5m;
- *greater-than-or-equal-to*—Same as *greater-than* but checks for greater than or equal to (>=)
- *less-than*
 - Returns: True or False
 - Syntax: less-than <LHS> <RHS> [time-range <range>]
 - Example: //Memory < 6000 MB in the last 5 minutes
 when less-than \$memory 6000 time-range 5m;
- *less-than-or-equal-to*—Same as *less-than* but checks for less than or equal to (<=)
- *equal-to*—Used for checking that one value is equal to another value.
 - Returns: True or False
 - Syntax: equal-to <LHS> <RHS> [time-range <range>]

- Example: `//Queue's buffer utilization % == 0`
`when equal-to $buffer-utilization 0;`
- *not-equal-to*—Same as *equal-to* but checks for negative condition (`!=`)
- *exists*—Used to check if some value exists without caring about the value itself. Meaning that some value should have been sent from the device.
 - Returns: True or False
 - Syntax: `exists <$var> [time-range <range>]`
 - Example: `//Has the device configuration changed?`
`when exists $netconf-data-change`
- *matches-with (for strings & regex)*—Used to check for matches on strings using Python regex operations. See [Python Regular Expressions](#) for details.

NOTE: LHS, or left hand side, is the string in which we are searching; RHS, or right hand side, is the match expression. Regular expressions can only be used in RHS.

- Returns: True or False
- Syntax: `matches-with <LHS> <RHS> [time-range <range>]`
- Example: `//Checks that ospf-neighbor-state has been UP for the past 10 minutes`
`when matches-with $ospf-neighbor-state "^UP$" time-range 10m;`
- *does-not-match-with (for strings & regex)*—Same as *matches-with* but checks for negative condition
- *range*—Checks whether a value, X, falls within a given range such as minimum and maximum (`min <= X <= max`)
 - Returns: True or False
 - Syntax: `range <$var> min <minimum value> max <maximum value> [time-range <range>]`
 - Example: `//Checks whether memory usage has been between 3000 MB and 6000 MB in the last 5 minutes`
`when range $mem min 3000 max 6000 time-range 5m;`
- *increasing-at-least-by-value*—Used to check whether values are increasing by at least the minimum acceptable rate compared to the previous value. An optional parameter that defines the minimum acceptable rate of increase can be provided. The minimum acceptable rate of increase defaults to 1 if not specified.

- Returns: True or False

- Syntax:

increasing-at-least-by-value <\$var> [increment <minimum value of increase between successive points>]

increasing-at-least-by-value <\$var> [increment <minimum value of increase between successive points>] time-range <range>

- Example: Checks that the ospf-tx-hello has been increasing steadily over the past 5 minutes.

```
when increasing-at-least-by-value $ospf-tx-hello increment 10 time-range 5m;
```

- *increasing-at-most-by-value*—Used to check whether values are increasing by no more than the maximum acceptable rate compared to the previous value. An optional parameter that defines the maximum acceptable rate of increase can be provided. The maximum acceptable rate of increase defaults to 1 if not specified.

- Returns: True or False

- Syntax:

increasing-at-most-by-value <\$var> [increment <maximum value of increase between successive points>]

increasing-at-most-by-value <\$var> [increment <maximum value of increase between successive points>] time-range <range>

- Example: Checks that the error rate has not increased by more than 5 in the past 5 minutes.

```
when increasing-at-most-by-value $error-count increment 5 time-range 5m;
```

- *increasing-at-least-by-rate*—Used for checking that rate of increase between successive values is at least given rate. Mandatory parameters include the value and time-unit, which together signify the minimum acceptable rate of increase.

- Returns: True or False

- Syntax:

This syntax compares current value against previous value ensuring that it increases at least by value rate.

increasing-at-least-by-rate <\$var> value <minimum value of increase between successive points> per <second|minute|hour|day|week|month|year> [time-range <range>]

This syntax compares current value against previous value ensuring that it increases at least by percentage rate

increasing-at-least-by-rate <\$var> percentage <percentage> per <second|minute|hour|day|week|month|year> [time-range <range>]

- **Example:** Checks that the ospf-tx-hello has been increasing strictly over the past five minutes.

```
when increasing-at-least-by-rate $ospf-tx-hello value 1 per second time-range 5m;
```

- *increasing-at-most-by-rate*—Similar to increasing-at-least-by-rate, except that this checks for decreasing rates.

Using these operators in the *when* clause, creates a function known as a user-defined condition. These functions should always return true or false.

If evaluation of a *term* results in a match, then the action specified in the *Then* clause is taken. By default, processing of terms stops at this point. You can alter this flow by enabling the **Evaluate next term** button at the bottom of the *Then* clause. This causes Paragon Insights to continue *term* processing to create more complex decision-making capabilities like when-this and this, then that.

The following is a list of pre-defined actions available for use in the *Then* section:

- next
- status

Tagging

Tagging allows you to insert fields, values, and keys into a Paragon Insights rule when certain conditions are met. See "[Paragon Insights Tagging Overview](#)" on page 570 for more information.

Rule Properties

The **Rule Properties** block allows you to specify metadata for a Paragon Insights rule, such as hardware dependencies, software dependencies, and version history. This data can be used for informational purposes or to verify whether or not a device is compatible with a Paragon Insights rule.

Pre/Post Actions

In Paragon Automation, you can use the **Pre/Post-Action** tab in the Rules page to execute tasks that are preconfigured in action engine workflows. Action engine workflows are used to perform tasks that you can execute as CLI commands, NETCONF commands, or as commands in executable files. See "[Action Engine Workflow Overview](#)" on page 899 for more information.

When you run playbook instances on device groups, Paragon Automation executes pre-action tasks at the start of playbook instantiation. Pre-action tasks execute device configurations in a device group or issue notifications about device status to other applications. There is no dependency between multiple

pre-action tasks and between the execution of pre-action tasks and rules. If you configure more than one pre-action task in a rule, Paragon Automation executes all pre-action tasks simultaneously.

Paragon Automation executes post-action tasks when you stop a playbook instance. Post-action tasks remove any additional configurations added to devices through the pre-action tasks. However, post-action task configuration is optional.

Both pre-action and post-action tasks have the execute-once option. By default, execute-once is disabled. If you enable the execute-once option, then Paragon Automation executes the tasks only once on a device in a device group. Execute-once is applicable in the following cases:

- When you run multiple instances of a playbook on the same device group.
- When you include the same rule with a set of pre-action or post-action tasks in different playbooks and run the playbooks on the same device group.

Paragon Automation also checks for and resolves duplication of pre-action and post-action tasks before executing them. Duplication occurs when you configure a specific pre-action or post-action tasks in many rules that are included in a playbook.

NOTE: When you upgrade Paragon Automation, the application does not execute pre-action and post-action tasks that you deploy before the upgrade.

You can use the following steps to configure pre-action or post-action tasks:

1. Configure Action Engine Workflows. See ["Manage Action Engine Workflows" on page 903](#) for more information.
2. Configure Pre-Action or Post-Action tasks depending on your use case.

See ["Pre-Action Tasks" on page 358](#) to configure pre-action tasks and see ["Post-Action Tasks" on page 358](#) to configure post-action tasks.
3. Create a playbook with rules that have pre or post-action tasks. See ["Create a Playbook Using the Paragon Insights GUI" on page 309](#) for more information.
4. Run an instance of the new playbook on device groups to execute pre-action tasks. See ["Manage Playbook Instances" on page 312](#) for more information.
5. Stop the instance of the new playbook to execute post-action tasks. See ["Manage Playbook Instances" on page 312](#) for more information.
6. Monitor the status of pre-action and post-action tasks. See ["About the Workflows Monitor Page" on page 900](#) for more information.

RELATED DOCUMENTATION

[Understand Paragon Insights Topics | 320](#)

[Understand Sensor Precedence | 364](#)

About the Rules Page

Paragon Insight's primary function is collecting and reacting to telemetry data from network devices. A *rule* contains all the details and instructions to define how to collect and handle the telemetry data.

Insights ships with a set of default rules, which can be seen on the **Configuration > Rules** page of the GUI, as well as in GitHub in the [Paragon Insights rules](#) repository. You can also create your own rules.

NOTE: Rules, on their own, don't actually do anything. To make use of rules you need to use them to ["Add a Predefined Playbook" on page 308](#).

You can use the Rules page to perform the following tasks:

- Add a pre-defined rule. See ["Add a Predefined Rule" on page 340](#).
- Create a new rule. See ["Configure a Custom Rule in Paragon Automation GUI" on page 343](#).
- Edit a rule in Paragon Insights GUI. See ["Edit, Clone, Delete, and Download Rules" on page 341](#).
- Configure multiple sensors per device. See ["Configure Multiple Sensors per Device" on page 362](#).
- Configure sensor precedence in Rule Properties through CLI. See ["Configure Sensor Precedence" on page 365](#).
- Configure offset time. See ["Configure Offset Time" on page 474](#).
- Configure Paragon Insights to send notifications of gray failures. See ["Configure Paragon Insights Notification for LSP Gray Failures" on page 359](#)

Add a Predefined Rule

Juniper has created a set of predefined rules that you can use to gather information from various Juniper components and the networks they reside in. You can add these rules to Paragon Insights at any time. After installation, many default predefined rules appear in the Rules page. Predefined rules cannot be changed or removed; however, you can clone any rule (predefined or user defined) simply by clicking the

CLONE button on the upper-right part of the rule definition. A cloned rule goes to the *external* topic and can be re-configured at will.

To upload additional predefined rules to Paragon Insights :

1. Using a browser, go to <https://github.com/Juniper/healthbot-rules> and download the pre-defined rule file to your system.
2. In the Paragon Insights GUI, click the **Configuration > Rules** icon in the Paragon Automation menu.
3. Click the **↑ Upload Rule Files** button.
4. Click the **Choose Files** button.
5. Navigate to the rule file and click **Open**.
6. Select one of the following options:

Upload	Upload the file and save the rule within the defined topic area but do not deploy the updated configuration. You can use this option when, for example, you are making several changes and want to deploy all your updates at the same time.
Upload & Deploy	Upload the file, save the rule within the defined topic area, and immediately deploy the configuration.

RELATED DOCUMENTATION

| [Manage Rules](#)

Edit, Clone, Delete, and Download Rules

The following procedures show how to edit, clone, delete, and download rules.

1. Click **Configuration > Rules** in the Paragon Automation menu.
The Rules page appears.
2. Click the name of the rule (listed under a topic) in the left side of the Rules page.
For example, **check-interface-status**.
3. You can perform the following tasks:
 - Edit a rule – Modify the necessary fields in **Sensors** tab, **Fields** tab, **Trigger** tab, and so on.

After you edit a rule, you can **Save** or **Save & Deploy** the rule configuration.

- **Save** – Save your configuration changes but do not deploy the updated configuration. You can use this option when, for example, you are making several changes and want to deploy all your updates at the same time later.

If you choose only to save the changes, you can either commit or roll back the changes later. For more information, see "[Commit or Roll Back Configuration Changes in Paragon Insights](#)" on [page 185](#).

- **Save & Deploy** – Save the rule configuration in the GUI and deploy the configuration on your production environment. The ingest starts collecting telemetry data based on the configuration changes.
- (Optional) Clone a rule – Click **Clone** to create a copy of the rule configuration.

You can rename and edit the configuration details of the cloned rule. A cloned rule appears in the *external* topic.

- (Optional) Delete a rule – Click **Delete** located to the right of the rule name. A deletion confirmation message appears where you can **Delete** or **Delete & Deploy** a rule configuration.

If you delete a rule using **Delete**, the rule configuration is removed in Paragon Automation Platform only during next deployment. Until then, the application continues to collect sensor data based on the fields defined in the deleted rule. You can also rollback your deletion, if you used **Delete**. For more information, see "[Commit or Roll Back Configuration Changes in Paragon Insights](#)" on [page 185](#).

If you delete a rule using **Delete & Deploy**, the rule configuration is removed from Paragon Automation and you cannot rollback the change.

NOTE: Before you delete a rule, ensure that you remove the rule in Playbook instance(s) where you added the rule.

- (Optional) Download a rule configuration – Click **Download** located to the right of the rule name. You can open or save the rule configuration tarball in the pop-up window.

You must have pre-installed an application to open or extract tar files to open the rule configuration files.

To save the configuration tarball, select **Save File** in the download pop-up window, navigate to a location on your local system, and click **Save**. You can optionally rename the file before saving it in your system.

Configure a Custom Rule in Paragon Automation GUI

IN THIS SECTION

- [Create a New Rule Using the Paragon Automation GUI | 343](#)
- [Rule Filtering | 345](#)
- [Sensors | 346](#)
- [Fields | 348](#)
- [Vectors | 350](#)
- [Variables | 352](#)
- [Functions | 353](#)
- [Triggers | 355](#)
- [Rule Properties | 357](#)
- [Pre-Action Tasks | 358](#)

Create a New Rule Using the Paragon Automation GUI

To create a new rule using the Paragon Automation GUI, you'll first fill general descriptive information about the rule and then navigate through several rule definition blocks in the Rules page to provide the specific configuration for the Paragon Automation rule.

To start creating a new Paragon Automation rule:

1. Click the **Configuration > Rules** icon in the left-navigation bar. A list of Paragon Automation rules is displayed along the left side of the Rules page.
2. Click the add rule button (+ **Add Rule**).
3. Enter general descriptive information about the rule using the following input parameters:

Parameter	Description
Rule	<p>For a new rule, this parameter is pre-populated with <i>external / user_rule_random_characters</i>, for example, <i>external / user_rule_2p0ghk</i>. The fields separated by the slash (/) represent Paragon Automation topic name and Paragon Automation rule name, respectively.</p> <p><i>external</i> is the topic name used for user-defined topics. For the Paragon Automation rules pre-defined by Juniper, Juniper has curated a set of pre-defined device component-based topic names. For more information about Paragon Automation topics, see "Understand Paragon Insights Topics" on page 320.</p> <p>Replace the <i>user_rule_random_characters</i> rule name with a name that appropriately represents the rule's description such as packets-in, packets-out, system_memory, etc.</p>
Rule frequency	(Network rule only) Specify how often data for the network rule is collected by Paragon Automation . This setting is overridden if the rule is included in a frequency profile that is applied to a network group.
Description	(Optional) Enter a detailed description for the rule.
Synopsis	(Optional) Enter a brief description for the rule. The synopsis is displayed when you hover over the rule name listed along the left side of the Rules page.
Field Aggregation Time Range	This optional value defines how often Paragon Automation aggregates the data received by the sensor. This helps reduce the number of data point entries in the time series database.

4. (Network rule only) If the new rule is a network rule, toggle the Network rule switch to the right.

5. Configure the rule definition blocks as needed.

Located directly below the Synopsis input parameter, you'll find links to the following rule definition blocks: **Sensors**, **Fields**, **Vectors**, **Variables**, **Functions**, **Triggers**, and **Rule Properties**. The following sections describe the input parameters for each of these rule definition blocks.

6. Select one of the following options to save the new rule:

Save	Save the rule within the defined topic area but do not deploy the updated configuration. You can use this option when, for example, you are making several changes and want to deploy all your updates at the same time.
Save & Deploy	Immediately deploy the configuration and save the rule within the defined topic area.

Rule Filtering

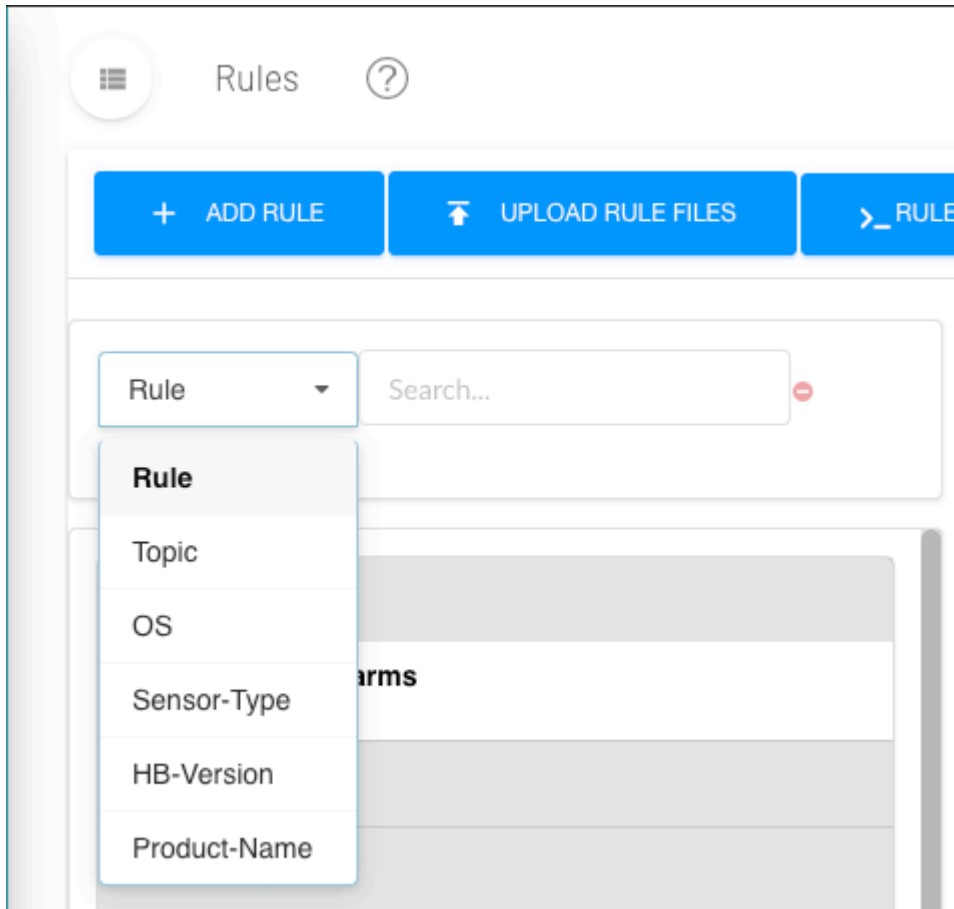
You can filter the Topics and Rules displayed on the left side of the Rules page. This allows you to quickly find rules that you are looking for. The search function works for topics, rules, sensor-types and other categories; working not only on titles, but also on the defined contents of rules.

The following procedure explains this filtering feature.

1. Navigate to **Configuration > Rules** in the left-navigation bar.

The **Rules** page is displayed. To the left of the rule definition area is a new section as shown in [Figure 18 on page 346](#) below.

Figure 18: Rule Filtering



2. From the pull-down menu, select the type of search you want to perform.
3. In the search field, begin entering your search text.

The topic list below shrinks to display only topics and rules that match your search criteria.

Sensors

Start configuring the new rule using the **Sensor** block. [Figure 19 on page 347](#) shows the sensor definition for the OpenConfig sensor `pppoe-error-statistics`.

Figure 19: A Sensor Definition

Rule: / Rule Frequency: Network Rule

Description: Collects the PPPoE error periodically and notifies in case of anomalies

Synopsis: Monitors PPPoE error statistics

Sensors Fields Vectors Variables Functions Triggers Rule Properties

+ Add sensor

pppoe-error-statistics

Sensor Name ⓘ

Sensor type

Sensor path *

Frequency *

1. Click the add sensor button (**+ Add Sensor**).

A new sensor definition appears and is named **Sensor_***random characters*, like Sensor_2kgf04.

You can configure more than one sensor for a rule.

2. Change the sensor name to something that makes sense for the rule you are defining.
3. From the drop-down list, choose the sensor type. You can choose one of: **OpenConfig**, **Native GPB**, **iAgent**, **SNMP**, **Syslog**, or **NetFlow**.

The required elements for defining the **Sensor Type** change depending on the selection you make. The frequency is expressed in #s, #m, #h, #d, #w, #y, or #o where # is a number and s, m, h, d, w, y specifies seconds, minutes, hours, days, weeks, years, and offset respectively. The o expression is used for defining an offset multiplier for use in formulas, references, triggers, learning periods, and hold-times.

The following list describes the elements that change based on your choice of **Sensor Type**. None of the other rule elements change because of a **Sensor Type** selection.

- **OpenConfig** **Sensor path** is defined from a drop-down list of available OpenConfig sensors. **Frequency** refers to how often, in seconds, the sensor reports to Paragon Automation . The frequency can be overridden if the sensor is included in a frequency profile.
- **Native GPB** **Sensor path** refers to the path for a Native GPB sensor. **Port** refers to the GPB port over which the sensor communicates with Paragon Automation .

- **iAgent** **File** is the name of a YAML-formatted file that defines the NETCONF-accessible sensor. **Table** is defined from a drop-down list of available PyEZ tables and views in the YAML file. **Frequency** refers to how often the sensor is polled by Paragon Automation and can be overridden by including the sensor in a frequency profile.

Based on the table you've selected, input fields for a target or dynamic arguments might also be provided. For these additional fields, you can do one of the following:

- Leave the input field blank. No default value will be applied.
 - Enter a fixed value that will remain constant.
 - Enter a variable name enclosed in double curly/flower brackets (for example, `{{test-variable}}`) The variable name must belong to a variable that was previously defined in the Paragon Automation rule, and the variable's **Type** option must be set to **Sensor Argument**.
- **SNMP** **Table** is defined from a drop-down list of available SNMP tables. **Frequency** refers to how often, in seconds, Paragon Automation polls the device for data and can be overridden by including the sensor in a frequency profile.
 - **Syslog** **Pattern set** is a user-configured element that includes one or more patterns (you configure a pattern for each event you want to monitor). The **Maximum hold period** is used in advanced cases and refers to the maximum time that the system will wait when correlating events using multiple patterns within a pattern set.

NOTE: The syslog sensor requires some pre-configuration. See [Syslog Ingest](#) for more details.

- **Flow** **Template Name** is a Juniper-supplied built-in list of NetFlow v9 and IPFIX templates.

Fields

A sensor will generally carry information such as all information about interfaces on the device, chassis related information, system process and memory related information, and so on. After you configure sensors, you must mention Fields that process sensor information for a particular need. For example, you can define fields to capture administrative or operational status of an interface or set traffic count threshold.

To add a field:

1. Click the **Fields** link.

The screen updates and shows the defined field objects.

2. Click the add field button (+ **Add Field**).
3. Replace the random field name with a name that make sense for the rule you are defining, such as **interface-name**, **configured-threshold**, etc.
4. (Optional) Add descriptive text for the new field.
5. Set the appropriate **Field Type**. The options for field type are: string, integer, float, and unsigned integer. String is the default field type.

You can also select unsigned integer as a field type. An unsigned integer is a data type that can contain values from 0 through 4,294,967,295.

6. (Optional) Toggle the **Add to rule key** switch.

The add rule to key switch tells Paragon Automation that this field should be indexed and searchable. For example, when you enable this switch, the field name will be listed on the Devices page under the **Keys** column.

7. Select the appropriate ingest type (Field source) from the pull-down menu.

The following list shows the options available for the **Ingest type (Field source)** menu.

- **Sensor**–Use this or another sensor definition.
 - **Path**–Follow this Open Config or Netconf path within the sensor definition to gather specific data like the names of the interfaces. For iAgent sensors the **Path** refers to the path defined in the YAML file.
 - **Where**–Filter the available data to gather information about a specific element within, like a specific interface. This field can reference the **Variables** defined elsewhere within the rule. When referencing variables, use moustache notation, enclosed in slashes, such as: **{{interface_name}}**.
 - **Zero suppression**–For some sensors associated with devices running Junos OS, such as Junos Telemetry Interface Open Config and native GPB sensors, no field data is sent from the sensor when the data's value is zero. Enable the zero suppression switch to set the field data value to zero whenever no field data is sent from the sensor.
 - **Data if missing**–Specify a value as the default data value whenever no data is sent from the sensor. The format of the specified value should match the defined field type (string, integer, or float). If the zero suppression switch is also enabled, then the specified data-if-missing value is ignored, and the default sensor data value is set to zero.
- **Reference**–A reference to a field or trigger value from another rule.
 - **Data if missing**–Specify a value as the default data value whenever no reference data is fetched. The format of the specified value should match the defined field type (string, integer, or float).

- **Constant**–Use a constant when referring to a **Variable** defined within the rule, whose value doesn't change, such as **IO_Drops_Threshold**. A constant can also be a string or numeric value that does not change and is not a reference to a variable.
 - **Constant value**–Use moustache notation to reference the variable like this: `{{IO_Drops_Threshold}}`.
 - **Formula**–Select the desired mathematical formula from the **Formula** pull-down menu.

8. (Optional) Set the **Field aggregation time-range**. Located above the Fields tab with the general rule parameters, this periodic aggregation setting helps to reduce the number of data points entered in the database. For example, if the sensor settings specify that Paragon Automation ingests data every 10 seconds, you could configure this setting to aggregate and record the relevant field data, say, every 60 seconds. Note that when using this setting, any field-specific time ranges must use the same value.

NOTE: You can add fields and keys to rules based on whether the incoming data meets user-defined conditions defined in tagging profiles. Tagging profiles are defined in Paragon Automation under **Administration > Ingest Settings** on the left navigation bar. See [Paragon Insights Tagging](#) for details.

Vectors

(Optional) Now that you have a sensor and fields defined for your rule, you can define vectors.

A vector is used when a single field has multiple values or when you receive a value from another field.

The syntax of a vector is:

```
vector <vector-name>{
  path [$field-1 $field-2 .. $field-n];
  filter <list of specific element(s) to filter out from vector>;
  append <list of specific element(s) to be added to vector>;
}
```

1. Click on the **Vectors** link. [Figure 20 on page 351](#) shows the Vectors block for a newly added vector.

Figure 20: Vectors Block

2. Click the add vector button (**+ Add Vector**)
3. Replace the random vector name with a name that makes sense for your rule.
4. Select an ingest type from the drop-down list. The additional input fields will vary depending on the selection you make.

For path:

Parameter	Description
List of fields	Select a field from the drop-down list. The list of fields is derived from all of the defined fields in this rule.
Time-range	Specify a time range from which the data should be collected. The time range is expressed in #s, #m, #h, #d, #w, #y where # is a number and s, m, h, d, w, y specifies seconds, minutes, hours, days, weeks, and years respectively. For example, enter 7 days as 7d.

For formula:

Parameter	Description
Vector name	(Unique formula type only) Select a vector name from the drop-down list. The list of vectors is derived from all of the defined vectors in this rule.

(Continued)

Parameter	Description
Formula Type	<p>Select a formula type from the drop-down list:</p> <p>unique Creates a vector with unique elements from another vector.</p> <p>and Compares two vectors and returns a vector with elements common to both vectors.</p> <p>or Compares two vectors and returns a vector with elements from both vectors.</p> <p>unless Compares two vectors and returns a vector with elements from the left vector but not the right vector.</p>
Left vector	Select a vector name from the drop-down list. The list of vectors is derived from all of the defined vectors in this rule.
Right vector	Select a vector name from the drop-down list. The list of vectors is derived from all of the defined vectors in this rule.

Variables

(Optional) The **Variables** block is where you define the parts of the sensor that you are interested in. For example, a rule that monitors interface throughput needs to have a way to identify specific interfaces from the list of available interfaces on a device. The field details are discussed below.

1. Click on the **Variables** tab.
2. Click the add variable button (+ **Add Variable**)
3. Replace the random **Variable name** with a variable name that makes sense for your rule, such as **pem-power-usage-threshold**.

BEST PRACTICE: The accepted convention within Juniper for naming of elements within Paragon Automation is to always start with a lower-case letter and use hyphens to separate words. Make sure that your variable names are unique, clearly named, and follow a recognizable pattern so that others can understand what the variable is for. Any abbreviations should be used consistently throughout.

4. Set an appropriate default value in the **Default value** field.

Default values vary depending on field type. Integer field types use numeric default values, while string field types use strings to set exact defaults and regular expressions that allow you to set the default from a list. Any default values set during rule definition can be overridden at apply-time at either the device or device group level.

5. Select the appropriate variable type from the **Type** drop-down list.

Available field types are: Integer, Floating Point, String, Boolean, Device, Device Group, and Unsigned Integer.

You can also select unsigned integer as a variable type. An unsigned integer is a data type that can contain values from 0 through 4,294,967,295.

Functions

(Optional) Define any needed functions.

The **Functions** block allows users to create functions in a python file and reference the methods that are available in that file. The python file must be created outside of Paragon Automation. You must know about the method names and any arguments because you will need those when defining the functions.

In Paragon Automation, you can use a Python user-defined function to return multiple values. The values are stored in multiple fields in the database.

For example, consider that you have a function `example_function.py` that has three return values. When you call the `example_function.py` in a rule, the first return value in the user-defined functions (UDF) is stored in the rule field that calls the function. You only need to configure return fields, such as `r2` and `r3`, for the remaining two return values. You can configure these fields for return values in the Return List of the Functions tab.

In the time-series database, the name of Return List fields are prefixed with the name of the rule field that uses the UDF. For example, `rule_field_name-r2`.

[Figure 21 on page 354](#) shows an example configuration in the **Functions** block.

Figure 21: The Functions Block

Sensors Fields Vectors Variables **Functions** Triggers Rule Properties

+ ADD FUNCTION

used_percentage DELETE USED_PERCENTAGE

Function name * ?

Path to Function * ? Method Name * ?

Description ?

Arguments

Name Mandatory ⊖

Name Mandatory ⊖

Return List

Name Type ⊖

+ ADD ARGUMENT + ADD RETURN LIST

To configure a function:

1. Click on the **Functions** link.
2. Click **+ Add Function**.
3. Enter a function name. For example, **used-percentage**.
4. In the **Path to function** field, enter the name of the python file that contains the functions. These files must be stored in the **/var/local/healthbot/input/hb-default** directory. The list is populated with all the Python (.py) files in that directory.
5. In the **Method Name** field, enter the name of the method as defined in the python file. For example, **used_percentage**.
6. (Optional) Enter a description for the function in the description box.
7. (Optional) For each argument that the python function can take, click **+ Add Argument**.
 Each time you click the add argument button, you'll need to enter the name of the argument and set the toggle switch as to whether the argument is mandatory or not. The default is that none of the arguments are mandatory.
8. (Optional) If there are more than one argument with a return value in the function, click **+Add Return List**.
9. Enter a name for the return value and the data type, such as an integer.

Triggers

A required element of rule definition that you'll need to set is the trigger element. [Figure 22 on page 355](#) shows the **Triggers** block for the `system.memory/check-system-memory` rule. The field details are discussed below.

Figure 22: The Triggers Block

The screenshot displays the 'Triggers' configuration page. At the top, there are navigation tabs: Sensors, Fields, Vectors, Variables, Functions, Triggers (selected), and Rule Properties. On the left, there is a '+ Add Trigger' button and a list of triggers, with 're-memory-buffer-utilization' selected. The main area shows the configuration for this trigger. The 'Trigger name' field contains 're-memory-buffer-utilization'. Below it is the 'Frequency' field with the placeholder 'Enter trigger frequency'. The 'WHEN' section contains a condition with four fields: 'Left operand' ('\$re-memory-buffer'), 'Operator' ('>='), 'Right operand' ('\$re-memory-buffer-high-threshold'), and 'Time range' ('5m'). There is an '+ Add Condition' button below the condition. The 'THEN' section includes a 'Color' dropdown set to red, a 'Message' text area containing 'Srouting-engine memory buffer utilization(\$re-memory-buffer) exceed high threshold(\$re-memory-buffer-high-threshold)', and an 'Evaluate next term' toggle. At the bottom, there is an information box with an 'i' icon and the text: 'Functions can be used as Trigger actions too, define them using the 'Functions' menu at the top.'

Setting up triggers involves creating terms that are used to set policy. If the terms within a trigger are matched, then some action is taken. Terms can evaluate the fields, functions, variables, etc that are defined within the rule against each other or in search of specific values. Terms are evaluated in order from the top of the term list to the bottom. If no match is found, then the next term (if any) is evaluated until a match is found or until the bottom of the term stack is reached.

1. Click on the **Triggers** link.
2. Click on the add trigger button (**+ Add Trigger**).
3. Replace the random trigger name with one that makes sense for the trigger you are defining, such as **foo-link-operation-state**. We recommend using a name that is very unique to the rule and trigger to avoid having the same trigger name across two or more rules.
4. (Optional) Enter a value in the **Frequency** field. This value tells Paragon Automation how often the field data and triggers should be queried and evaluated. If no entry is made here, then the sensor

frequency is applied for this value. The frequency entered here can be entered as a multiple or, an offset, of the sensor frequency such as 2o. For example, if the sensor frequency is 10s and the trigger frequency is 2o, then the trigger frequency would be 20s (2*10s).

5. Click the add term button (+ **Add Term**).

The Term area will expand and show an add condition button, (+ **Add Condition**) in the **When** section and **Color** and **Message** fields in the **Then** section.

6. To define a condition that the term will evaluate, click the + **Add Condition** button.

The **When** section expands to show **Left operand**, **Operator**, and **Time range** fields.

NOTE: Setting a condition is not required. If you want to guarantee that a **Term** takes a specific action, don't set a condition. This could be useful, for example, at the bottom of a term stack if you want some indication that none of the terms in your trigger matched.

7. Select values from the pull-down menus for each of these fields.

Depending on which **Operator** is chosen, a new field, **Right operand** may appear in between the **Operator** and **Time range** fields.

The left and right operand pull-down menus are populated with the fields and variables defined in the rule. The operator field determines what kind of comparison is done. The time range field allows the trigger to evaluate things such as if there were any dropped packets in the last minute.

8. (Optional) Set values for the **Color** and **Message** fields, and add **Action Engine** information in the **Then** section.

These fields are the action fields. If a match is made in the condition set within the same term, then whatever action you define here is taken. A color value of green, yellow, or red can be set. A message can also be set and is not dependent on whether any color is set.

You can link action workflows (Action Engine) to rules while setting up triggers. You can also add input arguments while linking action workflows. The **Action Engine** section is enabled only with a PIN-Advanced license. For more information, see ["Action Engine Workflow Overview" on page 899](#) and ["Paragon Insights Licensing Overview" on page 1005](#).

If color or message are set, a toggle button labeled **Evaluate next term** appears at the bottom of the **Then** section. The default value for this button is off (not active).

NOTE: If no match is made in the **When** section of a term, the **Then** section is ignored. If this happens, the next term down, if any, is evaluated.

If a match is made in the **When** section, then the actions in the **Then** section, if any, are taken and processing of terms stops unless the **Evaluate next term** button is set to on (active).

Setting the **Evaluate next term** button allows you to have Paragon Automation make more complex evaluations like 'if one condition and another condition are both true, then take this action'.

Rule Properties

(Optional) Specify metadata for your Paragon Automation rule in the **Rule Properties** block. Available options include:

Attributes	Description
Version	Enter the version of the Paragon Automation rule.
Contributor	Choose an option from the drop-down list.
Author	Specify a valid e-mail address.
Date	Choose a date from the pop-up calendar.
Supported Paragon Automation Version	Specify the earliest Paragon Automation release for which the rule is valid.
Supported Device > Juniper Devices	<p>Choose either Junos or Junos Evolved. Device metadata includes Product Name, Release Name, Release Support (drop-down list), and platform. You can add metadata for multiple devices, multiple products per device, and multiple releases per product.</p> <p>You can select default sensors that you want to apply to all supported devices. You also have the option to select default sensors that you want to apply to Juniper devices, and to Juniper devices that run a specific OS.</p>
Supported Device > Other Vendor Devices	<p>You can add vendor identifier, vendor name, product, platform, release, and operating system-related information for non-Juniper vendors.</p> <p>You can select default sensors that you want to apply to all non-Juniper devices.</p>
Helper Files	Specify files that are required by the Paragon Automation rule.

Pre-Action Tasks

IN THIS SECTION

- [Post-Action Tasks | 358](#)

Before you configure the pre-action tasks in rules, you must configure Action Engine workflows. See ["Manage Action Engine Workflows" on page 903](#) to configure Action Engine workflows.

To configure pre-action task:

1. Click the **Pre/Post Actions** tab on the Rules page.
2. In the Pre-Action section, select an action engine workflow in the Action Engine list.
3. Enter the input argument from the device list.

The list shows arguments you previously configured in the selected action engine workflow as the pre-action task.

4. Enable execute-once if you want Paragon Automation to execute the pre-action task only once on each device in a device group.

5. Do one of the following:

- Click **Save** to save your configuration changes but do not deploy the updated configuration. You can use this option when, for example, you are making several changes and want to deploy all your updates at the same time later.

See ["Commit or Roll Back Configuration Changes in Paragon Insights" on page 185](#) for more information.

- Click **Save and Deploy** to save the rule configuration in the GUI and deploy the configuration.

When you include the rule in a playbook and run a playbook instance on a device group, Paragon Automation executes the pre-action task while ingesting telemetry data.

You can monitor the status of the Action Engine Workflow embedded within a pre-action tab. See ["Manage Action Engine Workflows" on page 903](#) for more information.

Post-Action Tasks

Before you configure the post-action tasks in rules, you must configure Action Engine workflows. See ["Manage Action Engine Workflows" on page 903](#) to configure Action Engine workflows.

To configure post-action tasks:

1. Click the **Pre/Post Actions** tab on the Rules page.
2. In the Post-Action section, select an action engine workflow in the Action Engine list.
3. Enter the input argument from the device list.

The list shows arguments you previously configured in the selected action engine workflow as the post-action task.

4. Enable execute-once if you want Paragon Automation to execute the post-action task only once on each device in a device group.
5. Do one of the following:

- Click **Save** to save your configuration changes but do not deploy the updated configuration. You can use this option when, for example, you are making several changes and want to deploy all your updates at the same time later.

See ["Commit or Roll Back Configuration Changes in Paragon Insights" on page 185](#) for more information.

- Click **Save and Deploy** to save the rule configuration in the GUI and deploy the configuration.

When you stop the playbook instance (with this rule), Paragon Automation executes the post-action task after it stops the service for playbooks.

You can monitor the status of the Action Engine Workflow embedded within a post-action tab. See ["Manage Action Engine Workflows" on page 903](#) for more information.

RELATED DOCUMENTATION

| [Add a Predefined Rule | 340](#)

Configure Paragon Insights Notification for LSP Gray Failures

Paragon Insights can notify Paragon Pathfinder if there are gray failures in Label Switched Paths (LSPs) or nodes connecting the LSPs in an MPLS network. Gray failures or short-lived failures are mainly caused by interface-related issues such as link flaps, traffic exceeding bandwidth thresholds, link or interface errors, network delay, or packet loss and node-related issues such as routing daemon CPU utilization exceeding threshold, routing protocol issues, memory leaks, chassis temperature rising beyond a set threshold, or line card failures. When an LSP consists of a faulty link due to a gray failure,

Paragon Insights can be configured to notify Paragon Pathfinder which marks the link or node for maintenance and removes the maintenance mode status when the cause of gray failure is resolved.

Before you begin to configure functions, you must have completed configuring the **interface.statistics/check-interface-flaps** rule. The **interface.statistics/check-interface-flaps** rule periodically collects link flap counts and notifies Paragon Pathfinder when the link flap counts exceeds a pre-defined threshold.

The following checklist ensures that you have configured the **interface.statistics/check-interface-flaps** rule in Rules page:

- Selected sensor and configured sensor frequency in **interface.statistics/check-interface-flaps** rule in **Sensor** tab.
- Added a field to take flap count from sensor in **Fields** tab.
- Added a field to set flap threshold in **Fields** tab.
- Added a field to capture interface name in **Fields** tab.
- Added a Term to check if link flaps are increasing compared to threshold value for a given time range in **Triggers** tab.
- Added a Term to check if link is stable in **Triggers** tab.

Paragon Insights has a number of pre-defined rules that collect interface statistics such as flap count, administrative state of the interface, neighbor interface states, and so on. Let us take the example of how you can configure a function in Paragon Insights that sends Paragon Pathfinder a maintenance mode alert when flap count increases consistently.

After you configure the **interface.statistics/check-interface-flaps** rule, you can configure functions to generate or remove link maintenance mode in the **Functions** tab of the rule. The Python function catalog *notify_pathfinder.py* has pre-defined functions for generating and removing maintenance modes. You can make a reference to these functions and set arguments (parameters).

To add a link maintenance function:

1. Go to **Configuration > Rules** in the left navigation bar.
The Rules page appears.
2. Search for *interface* in the search box of Rules page.
This brings up all pre-defined rules related to interfaces. Check for **interface.statistics/check-interface-flaps** and click on the rule.
3. Click the **Functions** tab.
4. Click the **Add Function** button.
5. Enter a name for the function such as **generate_maintenance**.
6. Enter the description of the function in the **Description** field.
7. In the **Path to Function** field, enter **notify_pathfinder.py**.

8. In the **Method Name** field, enter `generate_link_maintenance`.

You can click the script icon near the field to open the catalog of functions available in `notify_pathfinder.py`. This shows a set of python functions with arguments (parameters). You must define the arguments mentioned in the function you choose.

9. Click the **Add Arguments** button to enter all the arguments specified in the `generate_link_maintenance` function that is listed in the `notify_pathfinder.py` catalog.

If you enable the *Mandatory* button, you must enter a value for the argument when you call the function.

10. Click **Save and Deploy**.

In the **Triggers** tab, you can call the *function name* you added and enter the value for the arguments (function parameters).

To call the generate maintenance mode function for link flaps in a trigger:

1. Click **Triggers** tab.
2. In the `link_flaps` trigger, click on the Term you configured to check if link flaps exceed the threshold limit in a given time.
3. Under the *Then* statement of the Term, click the **Add Function** button.
4. Select the *function name* you configured to generate maintenance link from the drop-down menu.

All the arguments (function parameters) you entered in the **Functions** tab appear here.

5. Enter `$` followed by argument name in the value field for each argument. You can also select a suitable input from the drop-down menu. For example, `$interface-name` for the argument `ifd_name`, `$mnt_time` for argument `mnt_time`, and so on.
6. Click **Save and Deploy**.

You can follow the same instructions to add `remove_link_maintenance` function in **Functions** tab and define argument values for that function under the Term in **Triggers** tab that checks if an interface is stable.

The four main functions that enable Insights to trigger notifications are `generate_link_maintenance` and `remove_link_maintenance` for any interface related issues and `generate_node_maintenance` and `remove_node_maintenance` for any node related issues.

To check nodes or interfaces that are under maintenance, navigate to **Network > Topology** in the left navigation bar. Click **ellipse (...)** in the panel with node, link and tunnel tabs. Select Maintenance in the drop-down menu. The Maintenances page displays details of links and nodes kept under maintenance.

RELATED DOCUMENTATION

[About the Topology Page | 677](#)

Configure Multiple Sensors per Device

Sp-admins or users with create access privilege must note the following guidelines when configuring multiple sensors:

- When adding multiple sensors to rules, you must ensure that there are no overlapping data or keysets received from the sensors applied to a device. Overlapping keysets can result in duplicate data points, overwriting of data points, and inaccurate evaluation of data. To avoid this, you can use filter expressions such as the *where* statement in **Fields**.
- When you add multiple sensors in a rule in Paragon Insights GUI, you must set a common value for all sensors in the following fields:
 - *Frequency* field in **Sensors** tab (sensor frequency)
 - *Field aggregation time-range* field in Rules page.
 - *Frequency* field in **Triggers** tab.
 - *Time range* field used in triggers, formula, and reference.

For example, when multiple sensors are added in a rule, all sensors applied to a device must have the same value for sensor frequency, irrespective of the type of sensors. If a rule has iAgent and OpenConfig sensors, the *Frequency* value in both sensors must be the same. This holds true for all the fields listed above.

NOTE: We recommend you use offset values if you cannot match the time range values on different sensors. For more information, see [Frequency Profiles and Offset Time](#).

However, the frequency you set in a frequency profile will override the frequency values set in multiple sensors in a rule.

- A Rule with multiple sensors is applied on all devices added in a particular device group. It is assumed that the devices in a device group support the types of sensors used in Paragon Insights rules.

Sometimes, not all devices in a device group can support the same type of sensor. For example, the device group DG1 has an MX2020 router with OpenConfig package installed and another MX2020 router configured without the OpenConfig package. The first MX2020 router would support OpenConfig sensor whereas, the second MX2020 router would not support the same sensor.

To avoid such scenarios, you must ensure that all devices in the same device group support the types of sensors used to collect information.

To add multiple sensors in a rule:

1. Navigate to **Configuration > Rules > Add Rule**.

Rules page appears.

2. Enter the Paragon Insights topic and rule name, rule description, synopsis, and optionally, the field aggregation time-range and rule frequency. For more information on these fields in Rules page, see [Paragon Insights Rules and Playbooks](#).

3. Click **Add Sensor** button in **Sensors** tab and fill in the necessary details based on the type of sensor you choose.

Repeat step 3 to add as many sensors as you need for your use case.

4. Configure fields for the sensors in the **Fields** tab.
5. Configure sensors in Rule Properties tab to set multiple sensors active.

You must enter all the sensors, earlier configured in the rule, under the supported-devices hierarchy in Rule Properties. For example, if you configured sensors s1, s2, and s3 in a rule, the Rule Properties configuration must also include the same sensors:

```
rule-properties {
    version 1;
    contributor juniper;
    supported-healthbot-version 4.0.0;
    supported-devices {
        sensors [s1 s2 s3];
    }
}
```

You can also write a new rule and upload it to the Paragon Insights GUI.

The rule must follow the curly brackets format ({ }) and indentation for hierarchical structure.

Terminating or leaf statements in the configuration hierarchy are displayed with a trailing semicolon (;) to define configuration details, such as supported version, sensors, and other configuration statements.

6. Click **Save & Deploy** to apply the new sensors in your network or click **Save** to save the configurations of the new sensors and deploy them later.

RELATED DOCUMENTATION

[Configure Sensor Precedence | 365](#)

[Understand Sensor Precedence | 364](#)

Understand Sensor Precedence

Paragon Insights allows you to add multiple sensors per rule that can be applied to all the devices in a device group. In earlier releases, you could add only one sensor per rule. Each sensor generates data in a field table. If you add the different sensors in multiple rules, it results in as many field tables as the number of rules. When you add multiple types of sensors (such as OpenConfig or Native GPB) in a single rule in Paragon Insights, data from the sensors is consolidated in a single field table that is simpler to export or to visualize. The GUI for multiple sensor will be implemented in subsequent releases.

The following scenarios illustrate use cases for multiple sensors in a rule:

- In Paragon Pathfinder, there can be different native sensors that provide non-overlapping counter details for Segment Routing (SR) and Resource Reservation Protocol (RSVP) label switched paths (LSP). If the field table needs to be combined for the data collected from the LSPs, multiple sensors can be made active for a device in the same rule.
- If you want to get data for *ge* interface using iAgent sensor and for *fe* interface using Native GPB sensor, then you could use multiple active sensor for a device. You need to ensure non-overlapping data in this case by using Field filtering expression to filter by the interface name. Instead of interfaces, an sp-admin or a user with create access privilege can consider any other key performance indicators too.

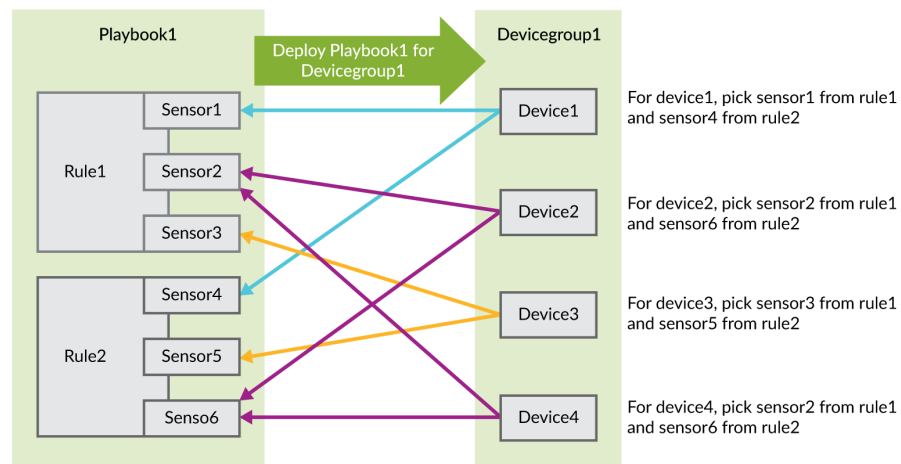
Paragon Insights also supports a feature that can be used with multiple sensors, called sensor precedence.

To make data collection from a sensor effective, devices in a device group must support a particular sensor as an ingest method. Select devices in a device group running an older version of operating system, devices from different vendors in a device group, or different products from the same vendor (such as EX, MX, and PTX routers from Juniper) are all scenarios that can cause challenges while applying a sensor to a device group. In such cases, you need to set a different sensor that is compatible with specific devices in a device group.

Paragon Insights enables you to set sensor precedence so that you can configure different sensors in each hierarchy of Rule Properties such as vendor name, operating system, product name, platform, and release version. This makes it possible to apply suitable sensors on multi-vendor devices in a heterogeneous device group. You can configure sensor precedence only through Paragon Insights CLI. The GUI for multiple sensor will be implemented in subsequent releases.

[Figure 23 on page 365](#) illustrates two rules each with multiple sensors. It is assumed that Rule Properties is configured for sensor precedence.

Figure 23: Representation of Sensor Precedence in Rules



Let us suppose Sensor1 in Rule 1 is OpenConfig and Sensor4 in Rule 2 is iAgent and Device1 runs Junos operating system (OS). If OpenConfig and iAgent were set as default sensors for Junos OS hierarchy in Rule Properties then, Device1 would receive data from Sensor1 and Sensor4 when the Playbook is deployed for the device group.

RELATED DOCUMENTATION

[Configure Multiple Sensors per Device | 362](#)

[Configure Sensor Precedence | 365](#)

Configure Sensor Precedence

IN THIS SECTION

- [Configure Sensor Precedence | 368](#)

In Paragon Automation Platform, the following fields in device configuration, required for sensor precedence, are extracted using base platform:

- *Vendor name*: Paragon Insights supports multiple vendors that include Juniper Networks, Cisco Systems, Arista Networks, and PaloAlto Networks.
- *Operating system*: Name of the operating system supported by vendors such as Junos, IOS XR, and so on.
- *Product*: Name of the family of products (devices) offered by a vendor. For example, MX routers, ACX routers, PTX routers from Juniper Networks.
- *Platform*: Particular member device in a series of products. For example, MX2020, ACX5400 and so on.
- *Release or version*: release version of the OS selected.

The configuration of above fields in the device hierarchy must match the sensor precedence specified in the Rule Properties. For example, if you include platform MX2020 in device configuration, the sensor precedence hierarchy must also include MX2020.

The following is a sample configuration of setting sensor precedence in Rule Properties:

```
rule-properties {
    version 1;
    contributor juniper;
    supported-healthbot-version 4.0.0;
    supported-devices {
        sensors interfaces-iagent;
        juniper {
            sensors interfaces-iagent;
            operating-system junos {
                products EX {
                    sensors interfaces-iagent;
                    platforms EX9200 {
                        sensors interfaces-iagent;
                        releases 17.3R1 {
                            sensors interfaces-iagent;
                            release-support min-supported-release;
                        }
                    }
                }
                platforms EX9100 {
                    sensors interfaces-oc;
                }
            }
            products MX {
                sensors interfaces-oc;
            }
        }
    }
}
```

```

    }
    products PTX {
        sensors interfaces-oc;
    }
    products QFX {
        sensors interfaces-oc;
    }
}
}
other-vendor cisco {
    vendor-name cisco;
    sensors interfaces-oc;
}
}
}

```

Sensor precedence mandates changes to the current hierarchy of configuration in Rule Properties. The hierarchy of the following elements has changed in Paragon Insights. The old hierarchy of the listed elements in Rule Properties is deprecated.

- *Releases*: The old statement hierarchy defined releases under product, and platform as a leaf statement under Releases. This hierarchy is deprecated.

```

products MX {
    releases 15.2R1 {
        release-support min-supported-release;    ### Deprecated
        platform All;                            ### Deprecated
    }
}

```

- *Operating system*: The old statement hierarchy defined operating system as a leaf statement for other vendors. This order is deprecated.

```

other-vendor cisco {
    vendor-name cisco;
    operating-system nexus;    ### Deprecated
    sensors [ s1 s2 ]; // Default sensors for cisco vendor
    operating-systems nxos {
        sensors [ s1 s2 ]; // Default sensors for cisco vendor with NX OS
        products NEXUS {

```

```
        sensors [ s1 s2 ]; // Default sensors for cisco vendor with NX OS and for
the specified product
        platforms 7000 {
            sensors [ s1 s2 ]; // Default sensors for cisco vendor with NX OS and
for the specified product and platform
            releases 15.8 {
                release-support only-on-this-release;
                sensors [ s1 s2 ]; // Sensors for cisco vendor with NX OS and for
the product, platform and version
            }
        }
    }
```

Configure Sensor Precedence

RELATED DOCUMENTATION

[Understand Sensor Precedence | 364](#)

[Configure Multiple Sensors per Device | 362](#)

Resources

IN THIS CHAPTER

- [Understand Root Cause Analysis | 369](#)
- [About the Resources Page | 372](#)
- [Add Resources for Root Cause Analysis | 375](#)
- [Configure Dependency Between Resources | 378](#)
- [Example Configuration: OSPF Resource Dependency | 383](#)
- [Edit Resources and Dependencies | 394](#)
- [Upload Resources | 395](#)
- [Download Resources | 396](#)
- [Clone Resources | 397](#)
- [Delete User-Generated Resources and Dependencies | 398](#)
- [Filter Resources | 399](#)

Understand Root Cause Analysis

IN THIS SECTION

- [Terms | 370](#)
- [Resources | 371](#)
- [Resource Dependencies | 371](#)
- [Use Cases | 372](#)

Paragon Automation monitors network resources — such as devices, interfaces, protocols, and label switched paths — through rules deployed in the network. Rules capture specific metrics called key performance indicators (KPI) for the network resources. Paragon Automation generates separate alarms

when different KPIs experience an anomaly or an error. These alarms do not establish a causal relationship between two error events.

To find the root cause of multiple errors, Paragon Automation must first link rules to their corresponding resource. For example, the rules that monitor temperature must be linked to the chassis resource or rules that monitor interface admin status must be linked to the interface resource. Users can link a rule to a resource when they configure a new resource or clone a system-generated resource. Secondly, Paragon Automation must perform root cause analysis by checking if an error in one resource leads to an error in another. To enable Paragon Automation to perform root cause analysis, you must configure dependency for the resources.

When Paragon Automation finds the root cause of several error events, it combines the alarms generated for these events into a top-down hierarchy based on the causal relationship between the events. Such a combined set of alarms is called a smart alarm.

Terms

The following list has frequently used terms and concepts connected with resources:

- *Resource*—A resource is a specific component that constitute the network.
For example, chassis, line card, protocols, system memory, interfaces, and so on.
- *Resource Instance*—Resources such as an interface, Flexible PIC Concentrators (FPC), router ids, or virtual private networks can have many instances. An instance is a specific realization of the resource. For example, an interface includes instances such as ge-0/0/1, et-1/0/0, xe-2/0/1, and so on.
- *Property*—Properties define unique characteristics of a resource. A property is comparable to fields in rules.
For example, neighbor-id and maximum transmission unit (MTU) are characteristics of routing-options resource and interface resource, respectively.
Neighbor id is a unique characteristic of protocol OSPF as a resource.
- *Key property in resource*— A key property uniquely identifies all instances of a resource.
For example, <variable>interface-name</variable> property uniquely identifies interface resource instance with name ge-0/0/0 from other instances. MTU cannot be a unique property.
- *Resource Dependency*—Defines the relationship between two resources.
- *Dependent resource*—In a resource dependency, a dependent resource is the one that depends on another resource. In a single-level resource dependency, a dependent resource is a child resource of another resource.

- *Dependency resource*—In a resource dependency, a dependency resource is the one that impacts another resource. In a single-level resource dependency, a dependent resource is a parent resource of another resource.

Resources

A resource can be part of a device or the network. Device resources can be an interface, Flexible PIC Concentrator (FPC), chassis, OSPF, etc. Network resources are resources that span multiple devices in a network, such as IPSec tunnels, VPN, etc.

As with rules, you configure resources under the *topic* hierarchy in Paragon Automation. The resource properties are derived from rules. When you configure interface as a resource, you can choose the specific interface rules from which Paragon Automation detects a particular interface resource property. The exact rules you select to identify a resource depend on your use case.

When you configure a resource property (such as interface name or MTU), you can refer multiple rules where the values configured for the property are different. A resource property aggregates instances from referenced rules.

Resource Dependencies

Resource dependency defines the relationship between a dependent resource (child resource) and a dependency resource (parent resource). While configuring dependency, you begin with a dependent resource and refer dependency resources that the child resource depends on. A dependency configuration also has terms that contain the logic to map dependency between two resources.

There are three types of dependency based on the type of resources involved, as described in [Table 68 on page 371](#). You can define dependencies between resources in the same device (Local Device and Network), between resources in different devices (Other Device), between a network resource to another network resource (Other Network), and a network resource to a device resource (Other Device).

Table 68: Types of Resource Dependency

Local Device and Network	Other Device	Other Network Group
Interface → line card	Interface 1 (device 1) → interface 2 (device 2)	VPN → interface 1 (device 1)
Line card → chassis	OSPF (device 1) → OSPF (device 2)	VPN → IPSec tunnel

You can configure multiple single-level dependencies for a resource. Consider the following chains of dependencies:

- VPN → LSP → interface → Line card

- VPN → interface → Line card

When you configure VPN as a resource, you can define VPN's dependency on Label Switched Paths (LSPs) as one dependency and VPN's dependency on interface as a second dependency. For LSP as a resource, you can define its dependency on interface. For interface as a resource, you can define its dependency on line card.

A dependency term logic can involve checking parts of a dependent (child) resource property with parts of dependency (parent) resource's property, checking all instances of the dependency (parent) resource, checking all devices that have dependency (parent) resource instances, checking all or specific device groups, and checking all or specific network groups. Paragon Automation supports *matches-with* and user-defined functions for dependency logic.

As dependency configurations can involve complex operations, Paragon Automation also allows you to execute such operations in a function. The function returns a Boolean value that can be used to check if a dependency exists between resources.

Use Cases

Resource and dependency configurations serve the following use cases:

- Root cause analysis — You can understand the root cause of a failure in your network. Refer the resource dependencies and use the relationship between resources to diagnose alerts generated by triggers in rules.
- Smart alarms — Through smart alarms, Paragon Automation links several resources that have a failure to another resource that is the cause of the failure.

RELATED DOCUMENTATION

| [About the Resources Page | 372](#)

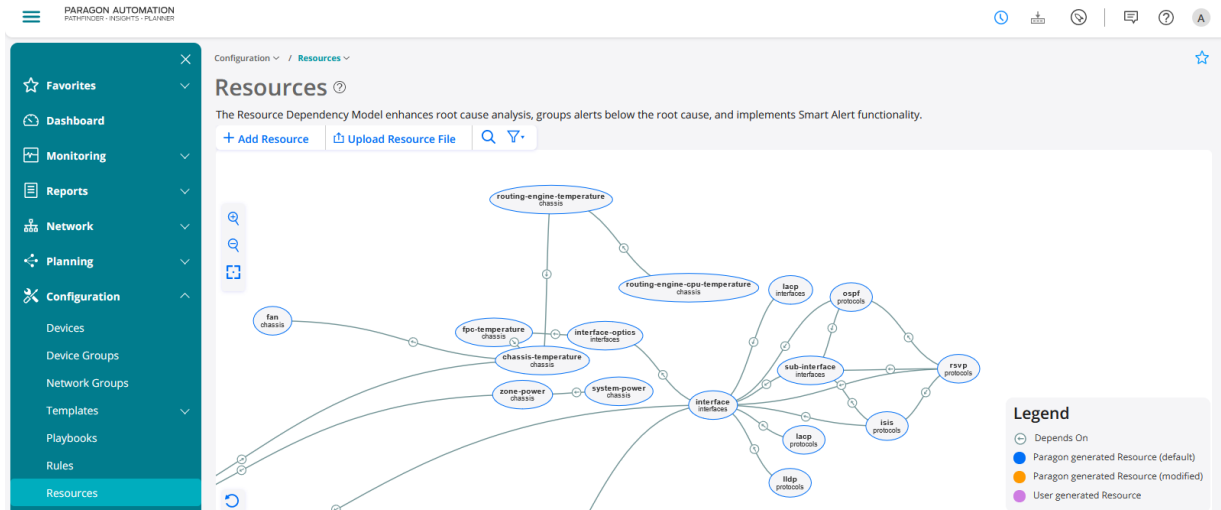
About the Resources Page

IN THIS SECTION

- [Tasks You Can Perform | 373](#)
- [Fields in Resources Table | 374](#)

You can access the Resources page from **Configuration>Resources**.

Figure 24: Resources and Dependency Visual Panel



The Resources page displays a visual representation of default resources in Paragon Automation. It also shows that a dependent (child) resource depends on the preceding dependency (parent) resource in the model. The visualization is in sync with Resources table so that you can work with resources from either the visual panel or the table. In Paragon Automation, you can reset the zoom level of the visual panel using the reset button after you zoom in or zoom out. You can also drag and change the position of resources on the visual panel. Paragon Automation retains any change to the position of resources, even after you leave the Resources page.

When you click a resource on the visual panel, Paragon Automation shows the key properties of the resource, the dependency resource (parent resource) and dependent resources (child resources) of the selected resource.

When you click the dependency arrow on the visual panel, Paragon Automation shows the terms and dependency type configured in the child resource.

Tasks You Can Perform

You can perform the following tasks in the Resources page:

- Add a resource from the resource visualization or table. See ["Add Resources for Root Cause Analysis" on page 375](#)
- Configure dependency resource. See ["Configure Dependency Between Resources" on page 378](#)
- Example Dependency Configuration of OSPF protocol. See ["Example Configuration: OSPF Resource Dependency" on page 383](#)

- Filter a resource from the resource visual panel or table. See ["Filter Resources "](#) on page 399
- Clone resources from the visual panel. See ["Clone Resources"](#) on page 397
- Upload resources on the visual panel or the resource table. See ["Upload Resources"](#) on page 395
- Download resources from the visual panel or the resource table. See ["Download Resources"](#) on page 396
- Edit a resource from the resource visualization or table. See ["Edit Resources and Dependencies"](#) on page 394
- Delete a resource from the resource visualization or table. See ["Delete User-Generated Resources and Dependencies"](#) on page 398

Fields in Resources Table

[Table 69 on page 374](#) displays the fields on the Resources page.

Table 69: Fields on the Resources Page

Field	Description
Name	Names of topics that expands to resources within the respective topic.
Description	Description you enter for the resource during resource configuration.
Keys	When you expand the topics, you can view key properties you configured for each resource.
Dependency	Parent resource that impacts the resource.
Generated by	Shows Paragon for system-generated (default) resources.

RELATED DOCUMENTATION

[Understand Root Cause Analysis](#) | 369

Add Resources for Root Cause Analysis

Resource configuration requires references to rules you want to link to the resource. Ensure that you configure the rules before you start resource configuration.

To configure a resource:

1. Select **Configuration > Resources**.

The Resources page appears.

2. Do one of the following:
 - a. Click **+Add Resource** on the visual panel.
 - b. Click **+** on the Resources table.

The Add Resource page appears.

3. Enter the details as described in [Table 70 on page 375](#).

4. Click **Save & Exit**.

The Save Resource Configuration dialog appears.

5. Do one of the following:
 - a. Click **Save and Deploy**.
You can save the resource configuration and deploy it.
 - b. Click **Save**.
You can only save the resource configuration but not deploy it.

If you save your configuration by only entering the details in the table, you complete the resource configuration. A message confirms the successful addition of the resource. You can click on the link in the confirmation message to go to the Alerts page where Paragon Automation lists smart alerts. See ["About the Alerts Page" on page 850](#) for more information.

You can see the new resource on the Resources page.

To connect this resource you add to another resource on the visual panel, you must configure dependency.

Table 70: Fields of Properties and Functions in Resource Configuration

Fields	Descriptions
General Information	
A resource, like rules, is defined under the <i>topic</i> hierarchy.	

Table 70: Fields of Properties and Functions in Resource Configuration (*Continued*)

Fields	Descriptions
Resource Name	<p>Enter name of the resource. The name can follow the regex pattern: [a-zA-Z][a-zA-Z0-9_-]* and can have up to 64 characters.</p> <p>For example, chassis, interface, or system.</p>
Topic	<p>Select the topic to which the resource belongs.</p> <p>For example, interface, chassis, or protocol.</p>
Description	<p>Enter a short description of the resource and the resources on which this resource is dependent.</p>
<p>Properties</p>	
<p>Properties are characteristics such as name, MTU or neighbor-id of a particular resource. A property of one resource can be matched with property of another resource to establish dependency.</p>	
Property Name	<p>Enter name of the resource property. The name can follow the regex pattern: [a-zA-Z][a-zA-Z0-9_-]* and can have up to 64 characters.</p> <p>For example, neighbor-id and interface name.</p>
Property Type	<p>Select the type of property input.</p> <p>You can choose from string, integer, floating point, or unsigned.</p>
Add as a Key	<p>If you enable a resource property as key, Paragon Automation uses this property to uniquely identify multiple instances of that property as a resource. You can mark more than one property as key.</p> <p>For example, in interface property that uses interface name as a key, Paragon Automation identifies ge-0/0/1, ge-1/2/0, ge-1/0/0 as unique instances of interface resource.</p>
<p>Source Configuration</p>	
<p>Rules and fields are the sources from which Paragon Automation discovers a resource property.</p>	

Table 70: Fields of Properties and Functions in Resource Configuration (*Continued*)

Fields	Descriptions
Rule Name	Add a rule name in the topic/rule format. For a resource property, you can add one or more rules as sources.
Field Name	Select a field configured in the rule.
(Optional) Functions You can enter a function in conditional statements that check for resource dependency.	
Function Name	Enter a function name follows the [a-zA-Z][a-zA-Z0-9_-]* pattern.
Path to Function	Select the file name where functions are defined. Paragon Automation supports only Python functions.
Method Name	Select the name of the function you want to execute from the given file.
Description	Enter a description of the function.
Arguments (Add Item)	
Name	Add argument or parameters defined in the function.
Mandatory	Toggle the switch on if the function arguments you added must be included when Paragon Automation calls the function.

RELATED DOCUMENTATION

[Understand Root Cause Analysis | 369](#)
[Configure Dependency Between Resources | 378](#)

Configure Dependency Between Resources

Before you configure dependency for a dependent (child) resource, you must complete configuring the dependency (parent) resources. See ["Add Resources for Root Cause Analysis" on page 375](#) for more information.

When you configure dependency for a dependent (child) resource, you configure terms that have the logical conditions to check for dependency. In Paragon Automation, you can change the order of dependency terms in the Dependency page. Paragon Automation executes the terms based on the sequence in which the terms appear in the GUI.

To configure a resource dependency:

1. Select **Configuration>Resource**.

The Resources page appears.

2. Select a resource from Resources table and click edit (pencil icon).

The Edit Resource page appears.

3. Click **Next** until you view the Dependency page.

4. Enter the details as described in [Table 71 on page 379](#).

5. Click **Save**.

The Dependency page appears.

6. Click **Next**.

You can see a visual representation of the resource and dependency configuration that you added.

7. (Optional) Click **View Tree** to view a collapsed view of the resource and the dependency configurations.

8. (Optional) Click **View JSON** to preview the JSON format of the configuration.

9. Click **Save & Exit**.

The Save Resource Configuration window appears.

- a. Click **Save** to only save the configuration.

Paragon Automation saves the configuration to form dependency between resource but does not generate smart alarms.

- b. Click **Save and Deploy** saves and deploys the configuration to form dependency between the resources and generate smart alarms.

You can see the new resource and its dependency mapped in the Resources page.

Table 71: Fields in Dependency Configuration

Fields	Description
<i>Dependent-resource-name depends on</i>	
Resource Name	Select the resource that impacts your dependent (child) resource.
<p>Add Variable</p> <p>Variables in this section are used to extract parts of the child resource's property. Such variables are used to check dependency using conditional statements.</p> <p>Variables configured here can be used to check dependency in all terms.</p>	
Name	<p>Enter a name that follows the [a-zA-Z][a-zA-Z0-9_-]* pattern.</p> <p>For example, interface-name-split.</p>
Expression	<p>Enter a regular expression to extract parts of a property value.</p> <p>See regex syntax for more information.</p> <p>For example, if you use the regular expression ".*(\d+)/(\d+)/\d+" on interface-name property of interface resource, you can extract line card number and PIC number.</p> <p>The example regex forms two capture groups to extract line card number and PIC number. The first capture group, which is line card number, can be referred as interface-name-split-1 and second capture group, which is PIC number, can be referred as interface-name-split-2.</p>
Resource	Select the name of the dependent (child) resource.
Field	Select a property in the child resource on which regex expression should be applied.

Table 71: Fields in Dependency Configuration (Continued)

Fields	Description
Case Sensitive	<p>Enable this field.</p> <p>If you enable this field, Paragon Automation ignores the case of resource properties when processing conditions.</p>
<p>Add Terms (Terms > Add Item)</p> <p>Terms contain the logic to check for a dependency relation between resources.</p>	
Term Name	<p>Enter a term name that follows the [a-zA-Z][a-zA-Z0-9_-]* pattern.</p>
Depends on Multiple Instances	<p>Enable this field if your dependency logic involves checking if a child resource property is dependent on multiple instances of a parent resource property.</p> <p>For example, you configured Aggregated Ethernet (link aggregation group) as a dependent (child) resource to interface resource. As AE depends on multiple links (interface resource), you must enable Depends on Multiple Instances knob.</p>
Dependency Type	<p>Select the type of dependency.</p> <p>Local Device and Network dependency is the default option.</p> <p>If you select Other Device dependency, configure details in For Every Device section.</p> <p>If you select Other Network dependency, configure details in For Every Network Group section.</p> <p>See "Understand Root Cause Analysis" on page 369 for more information on dependency types.</p>
Evaluate Next Term	<p>This field is disabled by default.</p> <p>Enable this field if your dependency logic involves checking condition in other terms.</p>

Table 71: Fields in Dependency Configuration (Continued)

Fields	Description
(Optional) Add Variable	<p>Enter the following details of a child resource property:</p> <ul style="list-style-type: none"> • Name • Regular expression • Resource name • Resource property field <p>You can configure a variable for a child resource property. Variable you configure here can only be used in conditional statements within the term.</p>
<p>For Every Device Selects a device from a list of devices in device groups.</p>	
Label As	Enter a name for the dependency (parent) resource instance that is selected. The name must follow the [a-zA-Z][a-zA-Z0-9_-]* pattern.
Across all device groups	Enable this field if you want Paragon Automation to check devices in all device groups for dependency.
In device groups	Enter the names of device groups. The child resource property will be checked against devices in the specified device groups for dependency.
<p>For Every Network Group Selects a device from a list of devices in network groups.</p>	
Label As	Enter a name for the dependency (parent) resource instance that is selected. The name must follow the [a-zA-Z][a-zA-Z0-9_-]* pattern.
Across all network groups	Enable this field if you want Paragon Automation to check all network groups for dependency.

Table 71: Fields in Dependency Configuration *(Continued)*

Fields	Description
In network groups	Enter the names of network groups. The child resource property will be checked against instances of network resource property in the specified network groups.
Locate Resource	
The locate resource iterates through all instances of a given resource.	
Resource Name	<p>Select a resource.</p> <p>If the resource type is Local Device & Network, the resource in the list is of the format <i>topic-name/resource-name</i>.</p> <p>If the resource type is Other Device or Other Network, then the resource in the list is of the format <i>for every device/network label-as:topic-name/resource-name</i>.</p>
Label As	<p>Enter a label name.</p> <p>Paragon Automation stores each instance of the selected resource in the label. For example, instances such as interface instances or OSPF sessions.</p>
(Optional) Add Variable	<p>Enter the following details of a dependency (parent) resource property:</p> <ul style="list-style-type: none"> • Name • Regular expression • Resource name • Resource property field <p>You can configure a variable for a property of the selected resource. Variable you configure here can only be used in conditional statements within the term.</p>

Table 71: Fields in Dependency Configuration (*Continued*)

Fields	Description
Conditions	<p>Conditions in Locate Resource are used to identify if the selected resource is correct one or not. If the condition does not match for a particular instance it picks the next instance and checks the condition. This continues until we get an instance where conditions matches, or all the instances of the resources are exhausted.</p> <p>To check dependency, select a resource property in left-hand side (LHS), a remote resource property in RHS and the operator to check for a condition.</p> <p>Paragon Automation supports matches-with as a condition. You can also add functions in the LHS field of the conditional statement.</p>

You can see the new resource and its dependency mapped in the Resources page.

RELATED DOCUMENTATION

[Understand Root Cause Analysis | 369](#)

[Add Resources for Root Cause Analysis | 375](#)

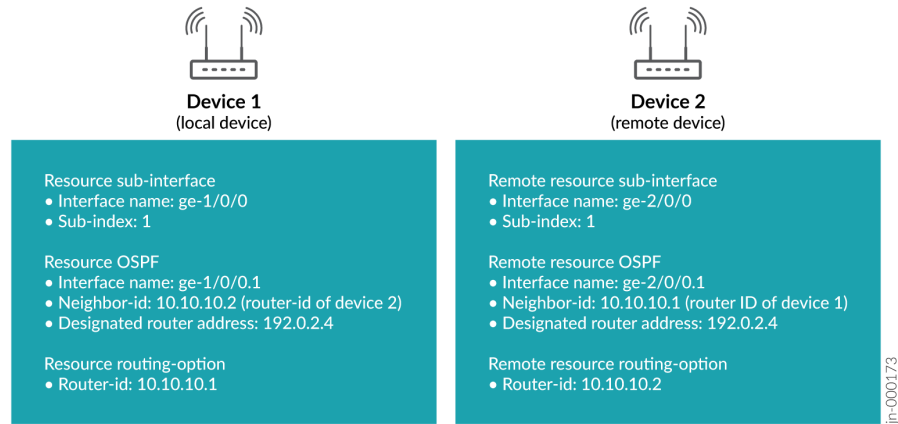
Example Configuration: OSPF Resource Dependency

Before you begin configuring the OSPF dependency, you must deploy the following configurations:

- Sub-interface as a resource with interface-name and sub-interface index as key properties.
- Routing options as a resource with router-id as a key property.

The following example configurations create OSPF protocol dependency between two devices. The OSPF protocol runs on an interface between two devices. So, the protocol forms two types of dependencies in the example configuration. The first dependency is between OSPF and the device, known as Local Device and Network dependency in the configuration. The second dependency is between OSPF protocol and the remote device, known as Other Devices dependency in the configuration. Paragon Automation establishes the two dependencies using the following properties in device 1 and device 2 as shown in [Figure 25 on page 384](#).

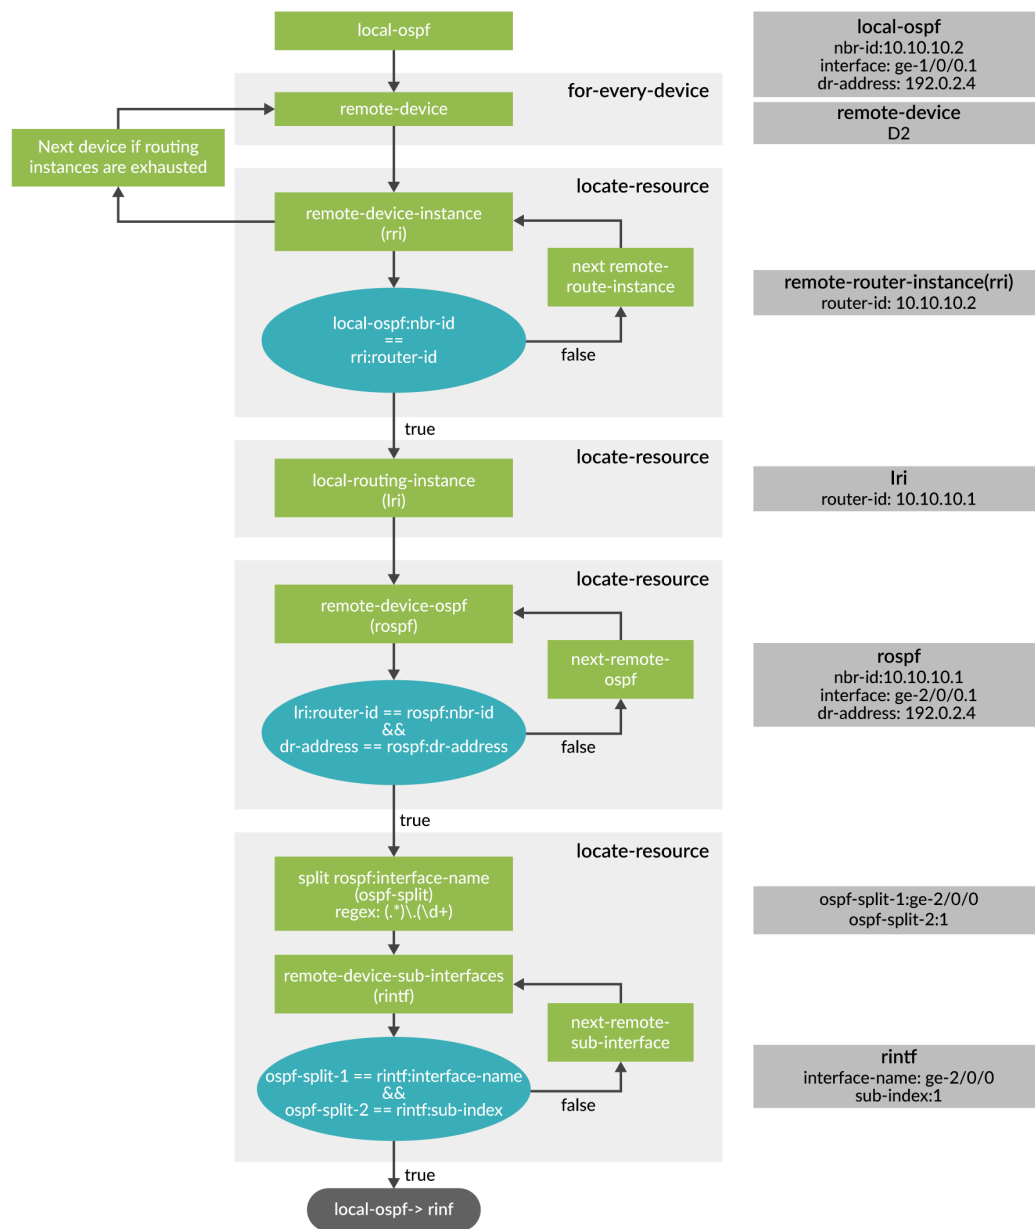
Figure 25: Properties Used to Establish OSPF Dependencies



The sub-interface properties and OSPF interface name property are used to create local dependency. Local dependency is established between a local device and the OSPF session at one end, and between the remote device and the OSPF session at the other end.

The neighbor ID, designated router address, and the device's router ID are used to create device-to-device (Other Device) dependency. [Figure 26 on page 385](#) shows how the iteration in locate resource configurations creates the device-to-device dependency.

Figure 26: Logic Flow of OSPF Other Device Dependency



To configure OSPF resource and dependencies:

1. Select **Configuration>Resource**.
The Resources page appears.
2. Select OSPF resource from the Resources table and click edit (pencil icon).
The Edit Resource page appears.
3. Enter the details as described in [Table 72 on page 386](#).

Table 72: Fields of Properties in OSPF Resource Configuration

Field	Description
General Information	
A resource, like rules, is defined under the <i>topic</i> hierarchy.	
Resource Name	<p>Enter ospf.</p> <p>The name must follow the regular expression pattern: <code>[a-zA-Z][a-zA-Z0-9_-]*</code> and can have up to 64 characters.</p> <p>For example, the name can also be Ospf-1.</p>
Topic	Select protocols .
Description	Enter a short description of OSPF dependencies.
Property Designated Router Address	
Properties are characteristics such as name, MTU, neighbor-id etc. of a particular resource. A property of one resource can be matched with the property of another resource to establish dependency.	
Property Name	<p>Enter dr-address.</p> <p>The name must follow the regex pattern: <code>[a-zA-Z][a-zA-Z0-9_-]*</code> and can have up to 64 characters.</p>
Property Type	Select String .
Add as a Key	<p>Not applicable.</p> <p>If you enable a resource property as key, Paragon Automation uses this property to uniquely identify multiple instances of that property as a resource. You can mark more than one property as key.</p> <p>For example, in interface property that uses the interface name as a key, Paragon Automation identifies <code>ge-0/0/1</code>, <code>ge-1/2/0</code>, <code>ge-1/0/0</code> as unique instances of the interface resource.</p>
Source Configuration	
Rules and fields are the sources from which Paragon Automation discovers a resource property.	

Table 72: Fields of Properties in OSPF Resource Configuration (Continued)

Field	Description
Rule Name	Select protocol.ospf/check-ospf-neighbor-information
Field Name	Select dr-address .
Property Interface Name	
Properties are characteristics such as name, MTU, neighbor-id etc. of a particular resource. A property of one resource can be matched with the property of another resource to establish dependency.	
Property Name	Enter interface-name . The name must follow the regular expression pattern: <code>[a-zA-Z][a-zA-Z0-9_-]*</code> and can have up to 64 characters.
Property Type	Select String .
Add as a Key	Enable this option. If you enable a resource property as key, Paragon Automation uses this property to uniquely identify multiple instances of that property as a resource. You can mark more than one property as key. For example, in interface property that uses the interface name as a key, Paragon Automation identifies <code>ge-0/0/1</code> , <code>ge-1/2/0</code> , <code>ge-1/0/0</code> as unique instances of the interface resource.
Source Configuration	
Rules and fields are the sources from which Paragon Automation discovers a resource property.	
Rule Name	Select protocol.ospf/check-ospf-neighbor-information
Field Name	Select interface-name .

Table 72: Fields of Properties in OSPF Resource Configuration (Continued)

Field	Description
Property Neighbor ID	
Properties are characteristics such as name, MTU, neighbor-id etc. of a particular resource. A property of one resource can be matched with the property of another resource to establish dependency.	
Property Name	Enter neighbor-id . The name must follow the regular expression pattern: <code>[a-zA-Z][a-zA-Z0-9_-]*</code> and can have up to 64 characters.
Property Type	Select String .
Add as a Key	Enable this option. If you enable a resource property as key, Paragon Automation uses this property to uniquely identify multiple instances of that property as a resource. You can mark more than one property as key. For example, in interface property that uses the interface name as a key, Paragon Automation identifies <code>ge-0/0/1</code> , <code>ge-1/2/0</code> , <code>ge-1/0/0</code> as unique instances of the interface resource.
Source Configuration	
Rules and fields are the sources from which Paragon Automation discovers a resource property.	
Rule Name	Select protocol.ospf/check-ospf-neighbor-information
Field Name	Select neighbor-id .

4. Click **Next** twice to go to the Dependency page.
The Dependency page appears.
5. Click **+** in the Dependency page.
A new Resource section appears.
6. Enter the details as described in [Table 73 on page 389](#).

Table 73: Local Device and Network Dependency Configuration

Fields	Description
OSPF depends on	
Resource Name	Select interfaces/sub-interfaces .
Add Terms (Terms>Add Item)	
Terms contain the logic to check for a dependency relation between resources.	
Term Name	Enter ifl-dependency . The name must follow the [a-zA-Z][a-zA-Z0-9_-]* pattern.
Depends on Multiple Instances	Enable this field. Paragon Automation checks with multiple instances of the property interface name and the property sub-interface index configured in resource sub-interface .
Dependency Type	Select Local Device & Network . See " Understand Root Cause Analysis " on page 369 for more information on dependency types.
Add Variables	Name: Enter ospf-interface-split . Expression: Enter (.*)\.(\\d+) Field: Enter \$interface-name .
Locate Resource	
Iterates over all instances of the interfaces and the sub-interface index in resource sub-interfaces to find local dependency.	
Resource Name	Select interfaces/sub-interfaces .
Label As	Enter ifl . Paragon Automation stores each instance of the interface resource to check for dependency.

Table 73: Local Device and Network Dependency Configuration (*Continued*)

Fields	Description
<p>Conditions</p> <p>Conditions in Locate Resource are used to identify if the selected resource is the correct one or not. If the condition does not match a particular instance, Paragon Automation picks the next instance and checks for the condition. This continues until we get an instance where the condition matches, or all the instances of the resource are exhausted.</p>	<p>The first condition checks for a match between the OSPF resource's interface property and the interface resource's interface instances.</p> <p>In left-hand side (LHS), select \$ospf-interface-split-1.</p> <p>Select matches-with as operator.</p> <p>In the right-hand side (RHS), select \$ifl:interface-name.</p> <p>Click + to add a second condition that checks for a match between the OSPF interface's sub-interface index and the resource sub interface's sub-interface index.</p> <p>In the LHS, select \$ospf-interface-split-2.</p> <p>Select matches-with as operator.</p> <p>In the RHS, select \$ifl:sub-interface-index.</p>

7. Click **Save**.

The Edit Resource page appears.

You can find the new term in Terms section. The OSPF's interface name and interface sub-index is checked against the sub interface resource's interface name and sub-index. When Paragon Automation finds a match, the interface from the OSPF resource monitored in a device forms a local dependency with the interface from the interface resource of the same device.

8. Click **+** in the Terms section.

You can add a second term to configure Other Device dependency for OSPF resource.

9. Enter the details as described in [Table 74 on page 390](#).

Table 74: Other Device Term Configuration

Fields	Description
Add Terms (Terms>Add Item)	

Table 74: Other Device Term Configuration (Continued)

Fields	Description
Term Name	Enter ifl-remote-dependency . A name must follow the [a-zA-Z][a-zA-Z0-9_-]* pattern.
Dependency Type	Select Other Device . See " Understand Root Cause Analysis " on page 369 for more information on dependency types.
For Every Device	
Selects a device from a list of devices in device groups.	
Label As	Enter remote-device The name must follow the [a-zA-Z][a-zA-Z0-9_-]* pattern.
Locate Resource	
Define condition to check if the local device's OSPF neighbor-id is the remote device's router-id.	
Resource Name	Select remote-device:protocols/routing-instance .
Label As	Enter remote-routing-instance
Conditions Conditions in Locate Resource are used to identify if the selected resource is the correct one or not. If the condition does not match a particular instance, Paragon Automation picks the next instance and checks for the condition. This continues until we get an instance where the condition matches, or all the instances of the resource are exhausted.	In the LHS, select \$neighbor-id In the operator field, select matches-with . In the RHS, select \$remote-routing-instance:router-id .
Locate Resource	
Use resource Routing Instance to collect router-id of the local device.	
Resource Name	Enter protocol/routing-instance .

Table 74: Other Device Term Configuration (Continued)

Fields	Description
Label As	Enter label name as local-routing-instance .
<p>Locate Resource</p> <p>Define condition to check if the remote device's OSPF neighbor-id matches with the router-id of the local device.</p>	
Resource Name	Enter remote-device: protocols/ospf .
Label As	Enter remote-ospf .
<p>Conditions</p> <p>Conditions in Locate Resource are used to identify if the selected resource is the correct one or not. If the condition does not match a particular instance, Paragon Automation picks the next instance and checks for the condition. This continues until we get an instance where the condition matches, or all the instances of the resource are exhausted.</p>	<p>In the LHS, enter \$local-routing-instance:router-id.</p> <p>In the operator field, select matches-with.</p> <p>In the RHS, enter \$remote-ospf:neighbor-id.</p> <p>Click + to add a second condition that checks for a match between the local device OSPF's designated router address and the remote OSPF's designated router address.</p> <p>In the LHS, enter \$dr-address.</p> <p>In the operator field, select matches-with.</p> <p>In the RHS, enter \$remote-ospf:\$dr-address.</p>
<p>Locate Resource</p> <p>Define a condition to check if the interface name and the sub-index in remote OSPF matches with the interface name and the sub-index of remote device's interface.</p>	
Resource Name	Enter remote-device: interfaces/sub-interface .
Label As	Enter remote-ifl .

Table 74: Other Device Term Configuration (Continued)

Fields	Description
<p>Conditions</p> <p>Conditions in Locate Resource are used to identify if the selected resource is the correct one or not. If the condition does not match a particular instance, Paragon Automation picks the next instance and checks for the condition. This continues until we get an instance where the condition matches, or all the instances of the resource are exhausted.</p>	<p>The first condition checks for a match between the remote OSPF resource's interface property and the remote interface resource's interface instances.</p> <p>In the LHS, enter \$ospf-remote-interface-split-1.</p> <p>In the operator field, select matches-with.</p> <p>In the RHS, enter \$remote-ifl:interface-name.</p> <p>Click + to add a second condition that checks for a match between the remote OSPF interface's sub-interface index and the remote interface instance's sub-interface index.</p> <p>In the LHS, enter \$ospf-remote-interface-split-2.</p> <p>In the operator field, select matches-with.</p> <p>In the RHS, enter \$remote-ifl:sub-interface-index.</p>

10. Click Next.

You can see a collapsed view of the resource and the dependency configuration you added.

11. (Optional) Click View JSON to preview the JSON format of the configuration.

12. Click Save & Exit.

The Save Resource Configuration window appears.

a. Click **Save** to only save the configuration.

Paragon Automation saves the configuration but does not generate smart alarms based on the new resource and dependency configuration you add.

b. Click **Save and Deploy** to deploy the saved configuration.

Paragon Automation saves and deploys the configuration to generate smart alarms.

You can see the OSPF resource and its dependency in the visual panel of the Resources page.

RELATED DOCUMENTATION

[Configure Dependency Between Resources | 378](#)

[Add Resources for Root Cause Analysis | 375](#)

Edit Resources and Dependencies

IN THIS SECTION

- [Edit a Resource | 394](#)
- [Edit Resource Dependency | 395](#)

Users can edit system resources, user-generated resources, and dependencies from the Resource Dependency Model page (**Configuration>Resources**). If you want to restore original configuration of system resources, you can restore by uploading the configuration files. Configuration files for system resources are available on Paragon Automation server and [GitHub](#). You can upload resource configuration from Rules page or Playbook page. See "[Add a Predefined Playbook](#)" on page 308 for more information.

NOTE: If you edit a system resource and later restore it to the original configuration, Paragon Automation continues to display the resource status as modified.

Edit a Resource

1. Select **Configuration>Resources**.
2. Do one of the following:
 - a. Select the resource from Resources table and click edit (pencil icon).
 - b. Select the resource on the resource model visualization and click **Edit** in the information pane of the resource.

The Edit Resources page appears.

3. Make changes to the configuration.

See "[Add Resources for Root Cause Analysis](#)" on page 375 for more information on resource configuration.

See "[Configure Dependency Between Resources](#)" on page 378 for more information on dependency fields.

4. Click **Save & Exit**.

The Save Resource Configuration window appears.

- a. Click **Save** to only save the configuration.

Paragon Automation saves the configuration but does not generate smart alerts based on the edits in resource configuration.

- b. Click **Save and Deploy** to save the configuration and deploy the configuration changes.

You can view changes in the information pane when you click on the edited resource.

SEE ALSO

| [Filter Resources](#) | [399](#)

Edit Resource Dependency

1. Select **Configuration>Resources**.

The Resource Dependency Model page appears.

2. Click on the arrow that shows dependency from the resource you want to edit to a parent resource.

The dependency page of the child resource appears.

3. Edit dependency configuration.

See "[Configure Dependency Between Resources](#)" on [page 378](#) for more information on dependency fields.

4. Click **Save & Exit**.

- a. Click **Save** to only save the configuration.

- b. Click **Save and Deploy** to save the configuration and deploy the configuration changes.

If you changed the dependency type or Term name, you can view the changes in dependency information when you hover over the arrow.

If you changed the depends on (parent) resource, you can view the edited resource connected to the new parent resource.

SEE ALSO

| [Delete User-Generated Resources and Dependencies](#) | [398](#)

Upload Resources

You can upload resource configuration from the Resources page **Configuration>Resources**. You can upload one resource file at a time but add multiple resource configurations in a resource file. Ensure that the .resource file is the file extension of a resource file.

If you add different resource configurations in a resource file, each configuration appears as a distinct resource on the Resources page. When you upload a resource configuration, the new configuration of that resource overrides the existing configuration if that resource exists in Paragon Automation.

NOTE: The name of each resource inside a resource configuration file (.resource file) must be unique.

To upload a resource file:

1. Click **Configuration>Resources**.

The Resources page appears.

2. On the visual panel, click **Upload Resource File**.

An Upload Resource File page appears.

3. Click **Browse** and select the resource file you saved in your local system.

4. Click **OK**.

You can see a confirmation message after the resource uploads successfully.

The new resource or resources appear on the Resources page in the visual panel.

RELATED DOCUMENTATION

[Clone Resources | 397](#)

[Download Resources | 396](#)

Download Resources

You can download resource configuration from **Configuration>Resources**.

To download a resource configuration:

1. Click the resource that you want to download.

A page that shows details of the resource appears.

2. On the page, click the download (down arrow) icon.

3. Click **Save File** and **OK** in the pop-up window to save the resource.

The resource configuration is downloaded to your local system as a compressed tar file.

You can change the resource configuration and re-upload the file with a different name using **Upload Resource File** in Resources page.

RELATED DOCUMENTATION

[Clone Resources](#) | 397

[Upload Resources](#) | 395

Clone Resources

To clone a resource:

1. Click the resource that you want to clone.

A page that shows details of the resource appears.

2. On the page, click the copy icon.

The Clone Resource page appears with the resource and dependency configurations of the resource you cloned.

3. Enter a name for the new resource in the Resource Name field.

4. Modify resource configuration details, if necessary.

See "[Add Resources for Root Cause Analysis](#)" on page 375 for more information.

5. Modify resource dependency, if necessary.

See "[Configure Dependency Between Resources](#)" on page 378 for more information.

6. Do one of the following:

- a. Click **Save and Deploy**.

Paragon Automation saves and deploys the resource and dependency configuration to generate smart alarms.

You can see the new resource on the Resources page.

- b. Click **Save**.

Paragon Automation saves the resource and dependency configuration. When you only save the configurations, Paragon Automation does not generate smart alarms.

You can see the new resource on the Resources page.

RELATED DOCUMENTATION

[Upload Resources | 395](#)

[Download Resources | 396](#)

Delete User-Generated Resources and Dependencies

IN THIS SECTION

- [Delete a Resource | 398](#)
- [Delete a Resource Dependency | 399](#)

NOTE: You cannot delete default (system-generated) resources.

Delete a Resource

1. Select **Configuration>Resources**.

The Resource Dependency Model page appears.

2. Do one of the following:

- a. Select the resource you want to delete on the resource visualization.

The information pane of the resource page appears.

- i. Click **Delete**.

- b. Select the resource from Resources table and click delete (trash icon).

A delete confirmation dialog appears.

3. Do one of the following:

- a. Click **Delete**.

Paragon Automation deletes the resource.

- b. Click **Delete and Deploy**.

Paragon Automation deletes the resource and deploy the configuration changes that occur as a result of the deletion.

SEE ALSO

| [Edit Resources and Dependencies](#) | 394

Delete a Resource Dependency

1. Select **Configuration>Resources**.
The Resource Dependency Model page appears.
2. Select the resource dependency you want to delete on the resource visualization.
The information pane of the dependency appears.
3. Click **Delete**.
A delete resource dependency confirmation dialog appears.
4. Click **OK**.
Paragon Automation removes the dependency you selected.

SEE ALSO

| [Edit Resources and Dependencies](#) | 394

Filter Resources

The Resources page shows the system resources. The visual panel expands as you add other resources. You can select resources and their dependencies through a keyword search and filter criteria in the Resources page.

When you filter or search a resource, the other resources gray out on the model visualization and the Resources table.

To filter resources:

1. Select **Configuration>Resources**.
The Resources page appears.
2. Click the filter icon (funnel) on the visual panel.
3. Click **Add Filter**.
The Add Criteria page appears.
4. Enter the details as described in [Table 75 on page 400](#).
You can see the filtered resources in the visual panel and the resources table.
5. (If you added multiple criteria) Do one of the following:
 - a. Select **AND** if you want Paragon Automation to filter resources based on more than one criteria.

- b. Select **OR** if you want Paragon Automation to filter resources based on either criterion.
- 6. Click **Save** if you want to reuse the filter you set.
The Save Filter page appears.
- 7. Enter a name for the filter.
- 8. Toggle the default switch on if you want the filter to appear first on the saved filter list.
- 9. Click **OK**.
- 10. Click the filter icon (funnel) icon to view all your saved filters.

If you click the filter icon (funnel) after saving filters, you can view only the saved filters and not the **Add Filter** menu. You must hide filters to view the **Add Filter** menu.

Table 75: Attributes in Add Criteria Page

Attribute	Description
Field	Select Topic Name, Resource Name, or Keys to filter the resources.
Condition	You can enter if the filter criterion is includes, equal to, or not equal for your input in the Value field.
Value	Enter a topic's name, resource's name, or a key name. The entry in the Value field depends on your selection in the Field.

RELATED DOCUMENTATION

| [Edit Resources and Dependencies](#) | 394

Graphs

IN THIS CHAPTER

- [Charts Overview | 401](#)
- [Manage the Charts Page | 404](#)
- [Monitor Network Devices from the Charts Page | 409](#)

Charts Overview

IN THIS SECTION

- [Types of Charts | 401](#)
- [Example Illustrations | 402](#)

Paragon Automation Platform offers three types of charts that include time series graphs, histograms, and heatmaps. You can use charts to monitor the status and health of your network devices. Charts allow you to visualize data collected by Paragon Insights from a device. To create charts, navigate to the **Monitoring > Graphs > Charts** page.

Types of Charts

You can create the following graphs from the **Charts** page.

Time Series Graphs

Time series graphs are charts that show data in a '2D' format where the x-axis indicates time and the y-axis indicates value. Time series graphs are useful for real-time monitoring, and also to show historical patterns or trends. This graph type does not provide insight into whether a value is 'good' or 'bad', but only shows 'the latest value'.

Histograms

Histograms are graphs that work quite differently than time series graphs. Rather than show a continuous stream of data based on when each value occurred, histograms aggregate data to show the distribution of values over time. This results in a graph that shows 'how many instances of each value'. Histograms also show data in a '2D' format, however, the x-axis indicates the value while the y-axis indicates the number of instances of the given value.


Heatmaps

Heatmaps are a combination of time series graphs and histogram elements. Heatmaps provide a '3D' view to help determine the deviations in data. Like a time series graph, the x-axis indicates time, and the y-axis indicates value. Heatmaps also include the number of instances of the given value similar to a histogram. Heatmaps also include a color dimension in the chart to add content. For each column in a heatmap, the bars indicate the various values that occurred. The color indicates how often the values occurred relative to the neighboring values. Within each vertical set of bars, the values that occurred more frequently show as 'hotter' with orange and red, while those values that occurred less frequently show as shades of green.

Example Illustrations

To help illustrate the three types of charts, consider the graphs shown below.

Figure 27:



All three graphs show the same data—the running 1-minute average of CPU utilization on a device over the last 24 hours. However, the way the data is depicted in each illustration is different.

- The time series graph provides the typical view; every minute it adds the latest data point to the end of the line graph. Time moves forward along the x-axis from left to right, and the data values are indicated on the y-axis. However, this graph doesn't show how often each data point has occurred.
- The histogram groups together the values to show *how many* of each data point there are. The tallest bar is the one between 30 and 40, which means that the most common 1-minute CPU average value is in the 30-40% range. And based on the y-axis, there have been over 350 instances of values in this range. The next most frequently occurring values are in the 40-50% range (almost 300 occurrences),

while the 0-10% range has almost no occurrences, suggesting that this CPU is rarely idle. This graph doesn't show how many instances of each data point occurred within a given time range.

- The heatmap makes use of elements from the time series graph and histogram. Each small bar indicates that some number of instances occurred within the value range shown in the y-axis, at the given time shown in the x-axis. The color indicates which value ranges, for each given time, occurred more than others. The vertical set of bars towards the right of the graph, at 18:00 illustrates this. In this example (at this zoom level), each column of vertical bars represents 12 minutes, and each small bar represents a bucket of 15 values. So the first (lowest) bar indicates that within this time range there were some values in the 0-14 range. The bar above indicates that within this time range there were some values in the 15-29 range, and so on. The color indicates which bars have more values than others. In this example, the third bar is red indicating that for those 12 minutes most of the values fell into the 30-44 range (in this example the count is 21). By contrast, the first bar is the most green indicating that for those 12 minutes the least number of values fell into the 0-14 range (in this example the count is 1). This 'heat' information is also supported by the histogram; the most frequently occurring values were those in the 30-40 range, which is the 'hotter' range in the heatmap.

RELATED DOCUMENTATION

[Manage the Charts Page | 404](#)

Manage the Charts Page

IN THIS SECTION

- [Add a Chart | 405](#)
- [Edit or Delete Charts | 408](#)

You can use charts to monitor the status and health of your network devices. Charts allow you to visualize data collected by Paragon Insights from a device, showing the results of rule processing. You can use the Charts page to add, edit, and delete a graph.

NOTE: Charts are refreshed every 60 seconds.

Add a Chart

The configuration model for charts is to create graph panels and group them into one or more canvases.

To create a new graph panel on a canvas:

1. Navigate to the **Monitoring > Graphs > Charts** page.

The Charts page appears.

2. Do any one of the following:

- To create a graph in a new canvas, click **+ New Canvas**.
- To create a graph within an existing canvas, select the desired canvas from the saved canvases and then click **Add Graph +**.

3. In the **General** section, provide the general descriptive information for the canvas (new canvas only) and graph:

Attribute	Description
Canvas Name	Enter a name for the canvas.
Description	(Optional) Enter a description for the canvas.
Graph Name	Enter a name for the graph.
Graph Type	Select the type of graph from the drop-down list. Options include Time Series , Histogram , and Heatmap .
Time Range	Choose the time range for the graph. In real-time graphs, the time range sets the x-axis range. For example, selecting 12 hrs means the x-axis shows the last 12 hours of data.

4. Move down to the Query section. In the From section, define from where the data for the graph is coming:

Attribute	Description
Group	Choose the device group or network group.
Device	Choose a device from the group.
Topic / Rule	<p>Choose the Paragon Insights topic/rule name.</p> <p>You can choose rolled-up measurements of fields you configured in rules. The rolled-up measurements are displayed in the list as <i>topic/rule_roll-up-order</i>. For example, my-rule_hourly and my-rule_yearly.</p> <p>The roll-up measurement's fields are listed as <i>original-field-name_aggregation-type</i>. For example, field-name_count and field-name_mean.</p> <p>Rolled-up measurements use aggregate functions on field data collected over a time to summarize the data into a single data point. You can configure how frequently the field data must be rolled-up or summarized. The default roll-up order (frequency) available in Paragon Automation are daily, hourly, monthly, weekly, or yearly.</p>

5. In the Select section, select the data field and apply aggregation and transformation types to the data:

Attribute	Description
Field	<p>Choose a field name.</p> <p>This list is derived based on the fields defined in the selected topic/rule.</p>
Aggregation	(Optional) In the drop-down list, choose a data aggregation type.
Transformation	(Optional) In the drop-down list, choose a data transformation type.

6. In the Where section, filter data based on field and KPI key:

Attribute	Description
Tag Key / Field	<p>(Optional) Choose a key or field. A key is an index field such as interface name.</p> <p>This list is derived based on the keys and fields defined in the selected topic/rule.</p> <p>You can set a maximum of two tags with WHERE conditions. You can use the <i>AND</i> option to satisfy both conditions and use <i>OR</i> option to satisfy either of the conditions you define.</p>
Condition	<p>Choose the conditional for the WHERE statement such as greater than, less than, matches with, equal to, and so on.</p> <p>The condition is used to check the right operand input against the left operand value.</p>

7. In the Group By section, specify how to group the data based on time interval, fill, and KPI keys:

Attribute	Description
\$_interval	<p>(Optional) Specify a time interval by which to group the data.</p> <p>If you set the time interval in GROUP BY field as \$_interval, then Grafana sets the time interval for data aggregation based on the number of points received in TSDB.</p>
fill(null)	<p>(Optional) Choose how to show a time interval when no data arrives from the device:</p> <p>null (default) Report the timestamp and null as the output value.</p> <p>none Report no timestamp and no output value.</p> <p>0 Report 0 as the output value.</p> <p>previous Report the value from the previous time interval as the output value.</p> <p>linear Report the results of linear interpolation as the output value.</p>
Tag Key (the + icon)	<p>(Optional) Choose a tag by which to group the data. A key is an index field such as interface name.</p> <p>This list is derived based on the keys and fields defined in the selected topic/rule.</p>

8. (Optional) Move down to the Visualization section, and define y-axis details.

You can enter the minimum value, maximum value, and label for y-axis. You can also choose a unit for y-axis such as bytes, Kilobytes, parts per million, and so on.

9. Click **Save** to save the graph and display the data.

TIP: You can use the following tips when working with graphs:

- To sort canvases on the Saved Canvas page, click on the column headings.
- To reorganize graphs on the screen, hover your mouse near the upper-left corner of a graph panel and click-and-drag it to the desired position.
- To resize a graph, hover your mouse over the lower-right corner of the graph panel and click-and-drag it to the desired size.
- To change the color of graph elements, click the color bar for the desired line item under the graph.
- To zoom in on a graph, click and drag across the desired section of the graph; to zoom out, double-click on the graph.
- To isolate an element on the graph, click its related line item under the graph; to view all elements again, click the same line item.

Edit or Delete Charts

In the Graph canvas, you can edit or delete a chart or delete a canvas.

- To edit a chart, click the pencil icon located in the top right corner of the chart itself.
- To delete a chart, click the trash can icon located in the top right corner of the chart itself.
- To delete a canvas, click the trash can icon located in the top right corner of the canvas.

RELATED DOCUMENTATION

| [Charts Overview](#) | 401

Monitor Network Devices from the Charts Page

SUMMARY

Read the following example configurations to understand the various ways you can monitor network devices from the **Charts** page.

IN THIS SECTION

- [Monitor Interface Flaps for a Single Interface | 409](#)
- [Monitor Interface Flaps for 'ge' Interfaces on a Device | 411](#)
- [Monitor System Memory Usage of all Devices in a Device Group | 413](#)
- [Monitor RE CPU Usage for Multiple Devices | 416](#)
- [Monitor RE CPU Usage for Multiple Devices side-by-side | 419](#)

You can use charts to monitor the status and health of your network devices. Charts allow you to visualize data collected by Paragon Insights from a device. For more information on creating charts and viewing charts, see ["Charts Overview" on page 401](#).

Monitor Interface Flaps for a Single Interface

The following example configuration describes how you can monitor interface flaps for a single interface.

1. Playbook used: **interface-kpis-playbook**
2. Example graph configuration:

Edit Graph

General

Canvas Name* TEST-canvas3 Description

Graph Name* vMX1-ge-003-interface-flaps Time Series 3 Hour

Query

FROM vMX vMX1 interface.statistics/check-interface-flaps

SELECT flaps mean() Transformation

WHERE interface-name = ge-0/0/3 AND

GROUP BY \$_interval fill(null) +

Visualization

Cancel Save

The example shows a configuration for vMX1 device in a canvas titled **Test-canvas3**. The query section takes interface flap count (flaps) as field input and computes the mean for flaps in **ge-0/0/3** interface. The time series graph in the configuration is captured over 3 hours.

- In the **Query** section, the **FROM** fields show the input for device-group name, device name, and rule, respectively.
- The **SELECT** fields contain inputs for a KPI key based on the rule you set earlier. If the rule you selected is `interface.statistics/check-interface-flaps`, the **SELECT** field could show options for flap count, flap threshold, and so on.
- The **Aggregation** field describes data aggregation computation to be performed on the KPI field data (flap counts in this example) such as mean, median, sum, and so on.
- The **Transformation** field (optional) describes computations performed on subsequent values in a single field. For example, rate of change, elapsed time, differences, and so on.
- The **WHERE** section allows you to set a condition to filter the field data. You must enter a left operand, a right operand, and a conditional operand such as **equal to**, **matches with**, **greater than**, **less than**, and so on. The left operand can be either a key or a field name.

In the example configuration, the condition defined checks if the interface-name is *equal to* ge-0/0/3.


- The **GROUP BY** section allows you to group the filtered data by a time interval you specify and handle output value in cases where no field data is available. You can choose group by time interval from the drop-down list.

If you set the time interval in **GROUP BY** field as **\$_interval**, then Grafana sets the time interval for data aggregation based on the number of points received in the time series database (TSDB).

- In the **Visualization** section, you can enter the y-axis minimum and maximum values, y-axis label and choose a unit such as bytes, kilobytes, and so on.

In this example configuration, the **Visualization** section is not configured.

3. Sample graph panel output:



The graph shows interface flaps for **ge/0/0/3** interface in device vMX1.

Monitor Interface Flaps for 'ge' Interfaces on a Device

The following example configuration describes how you can monitor interface flaps for all 'ge' interfaces on a device in a single graph.

1. Playbook used: **interface-kpis-playbook**

2. Example graph configuration:

Edit Graph

General

Canvas Name* TEST-canvas3 Description

Graph Name* vMX--interface-flaps Time Series 3 Hour

Query

FROM vMX vMX1 interface.statistics/check-interface-flaps

SELECT flaps mean() Transformation

WHERE interface-name =~ (matches with) ge AND

GROUP BY \$_interval fill(null) interface-name +

Visualization

Cancel Save

The example shows the configuration for vMX1 device in a canvas titled **Test-canvas3**. The query section takes interface flap count (flaps) as field input and computes the mean for flaps in all **ge** interfaces. The time series graph in the configuration is captured over 3 hours.

- In the **Query** section, the **FROM** fields shows the input for device-group name, device name, and rule, respectively.
- The **SELECT** fields contain inputs for a KPI key based on the rule that you set earlier. If the rule you selected is `interface.statistics/check-interface-flaps`, the **SELECT** field could show options for flap count, flap threshold, and so on.
- The **Aggregation** field describes data aggregation computation to be performed on the KPI field data (flap counts in this example) such as mean, median, sum, and so on.
- The **Transformation** field (optional) describes computations performed on subsequent values in a single field. For example, rate of change, elapsed time, differences, and so on.
- The **WHERE** section allows you to set a condition to filter the field data. You must enter a left operand, a right operand and a conditional operand such as **equal to**, **matches with**, **greater than**, **less than**, and so on. The left operand can be either a key or a field name.

In the example configuration, the condition defined checks if the interface-name *matches with* `ge`.


- The **GROUP BY** section allows you to group the filtered data by a time interval you specify and handle output value in cases where no field data is available. You can choose group by time interval from the drop-down list.

If you set the time interval in **GROUP BY** field as **\$_interval**, then Grafana sets the time interval for data aggregation based on the number of points received in TSDB.

- In the **Visualization** section, you can enter the y-axis minimum and maximum values, y-axis label and choose unit such as bytes, kilobytes, and so on.

In this example configuration, the **Visualization** section is not configured.

3. Sample graph panel output:



The graph shows flap counts of three ge interfaces — **ge-0/0/0**, **ge0/0/1**, and **ge-0/0/2** in the vMX1 device.

Monitor System Memory Usage of all Devices in a Device Group

The following example configuration describes how you can monitor system memory usage for all devices in a device group in a single graph.

1. Playbook used: **system-kpis-playbook**

2. Example graph configuration:

Edit Graph

General

Canvas Name* Description

Graph Name* Time Series

Query

FROM

SELECT

WHERE

GROUP BY

FROM

SELECT

WHERE

GROUP BY

FROM

SELECT

WHERE

GROUP BY

Visualization

Cancel Save

This example configuration shows the query section configured for three devices – vMX1, vMX2, and MX240 – in a single canvas titled **Canvas1**. The query section takes the routing engine's system memory buffer (re-memory-buffer) as field input, computes mean for a time interval set by the user or Grafana. The example configuration is set for the time series graph captured over 6 hours.

- In the **Query** section, the **FROM** fields show the input for device-group name, device name, and rule, respectively.
- The **SELECT** fields contain inputs for a KPI key based on the rule you set earlier. If the rule you selected is `system.memory/check-system-memory`, the **SELECT** field could show options for system memory buffer value, system memory high threshold, system memory low threshold, and so on.
- The **Aggregation** field describes data aggregation computation to be performed on the KPI field data (system memory buffer in this example) such as mean, median, sum, and so on.
- The **Transformation** field (optional) describes computations performed on subsequent values in a single field. For example, rate of change, elapsed time, differences, and so on.
- The **WHERE** section allows you to set a condition to filter the field data. You must enter a left operand, a right operand and a conditional operand such as **equal to**, **matches with**, **greater than**, **less than**, and so on. The left operand can be either a key or a field name.

In this example, no conditions are set in the **WHERE** section.

- The **GROUP BY** section allows you to group the filtered data by a time interval you specify and handle output value in cases where no field data is available. You can choose group by time interval from the drop-down list.

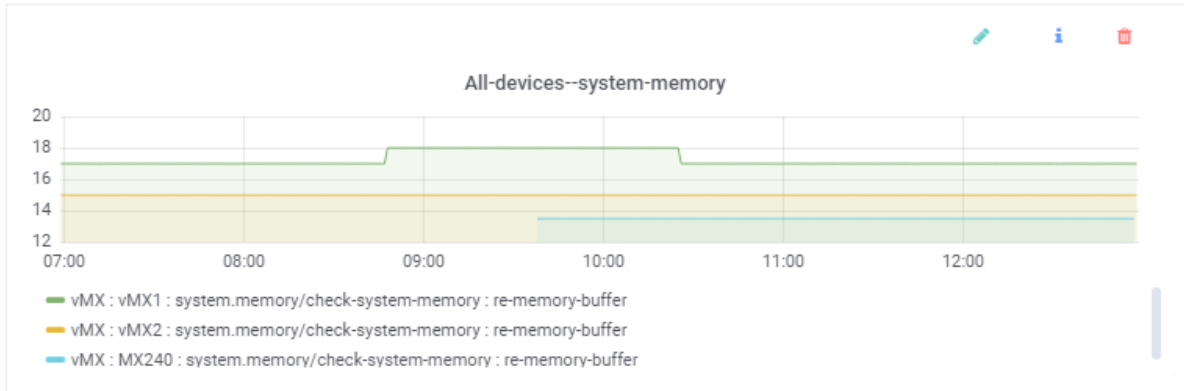
If you set the time interval in **GROUP BY** field as **\$_interval**, then Grafana sets the time interval for data aggregation based on the number of points received in TSDB.

In this example configuration, the group by interval is set at 1 minute for vMX1 and vMX2 devices whereas, Grafana would decide the time interval to group by aggregated data for MX240 device.

- In the **Visualization** section, you can enter the y-axis minimum and maximum values, y-axis label and choose unit such as bytes, kilobytes, and so on.

In this example configuration, the **Visualization** section is not configured.

3. Sample graph panel output:



The graph shows mean system memory buffer values aggregated per minute for vMX1 and vMX2, and for an unknown time interval for MX240 device.

Monitor RE CPU Usage for Multiple Devices

The following example configuration describes how you can monitor RE CPU usage for multiple devices in a single graph.

1. Playbook used: **system-kpis-playbook**
2. Example graph configuration:

Edit Graph

General

Canvas Name* Description

Graph Name* Time Series 6 Hour

Query

FROM

SELECT

WHERE

GROUP BY

FROM

SELECT

WHERE

GROUP BY

FROM

SELECT

WHERE

GROUP BY

Visualization

Y-Axis

This example configuration shows the query section and visualization section configured for three devices – vMX1, vMX2, and MX240 – in a single canvas titled **Canvas1**. The query section takes the routing engine’s CPU utilization (re-cpu-utilization) as field input, computes the mean every minute for the three devices, and visualizes the mean CPU utilization per minute as percentage. The example configuration is set for the time series graph captured over 6 hours.

- In the **Query** section, the **FROM** fields shows the input for device-group name, device name, and rule, respectively.
- The **SELECT** fields contain inputs for a KPI key based on the rule that you set earlier. If the rule you selected is `system.cpu/check-system-cpu`, the **SELECT** field could show options for CPU utilization value, CPU utilization high threshold, CPU utilization low threshold, and so on.
- The **Aggregation** field describes data aggregation computation to be performed on the KPI field data (CPU utilization in this example) such as mean, median, sum, and so on.
- The **Transformation** field (optional) describes computations performed on subsequent values in a single field. For example, rate of change, elapsed time, differences, and so on.
- The **WHERE** section allows you to set a condition to filter the field data. You must enter a left operand, a right operand and a conditional operand such as **equal to**, **matches with**, **greater than**, **less than**, and so on. The left operand can be either a key or a field name.

In this example, no conditions are set in the **WHERE** section.

- The **GROUP BY** section allows you to group the filtered data by a time interval you specify and handle output value in cases where no field data is available. You can choose group by time interval from the drop-down list.


If you set the time interval in **GROUP BY** field as `$_interval`, then Grafana sets the time interval for data aggregation based on the number of points received in TSDB.

In this example configuration, the group by interval is set at 1 minute for vMX1, vMX2, and MX240 devices.

- In the **Visualization** section, you can enter the y-axis minimum and maximum values, y-axis label and choose unit such as bytes, kilobytes, and so on.

In this example configuration, the y-axis range for mean CPU utilization per minute is set on a scale of 0 to 100 percentage.

3. Sample graph panel output:



The graph displays mean CPU utilization in percentage for vMX1, vMX2, and MX240 devices.

Monitor RE CPU Usage for Multiple Devices side-by-side

The following example configuration describes how you can monitor RE CPU usage for multiple devices side-by-side.

1. Playbook used: **system-kpis-playbook**
2. Example graph configurations:

Add Graph

General

Canvas Name* Description

Graph Name* Time Series 6 Hour

Query

FROM

SELECT

WHERE

GROUP BY

Visualization

Edit Graph

General

Canvas Name* Canvas1 Description

Graph Name* MX240-CPU-Utilization Time Series 6 Hour

Query

FROM vMX MX240 system.cpu/check-system-cpu

SELECT re-cpu-utilization mean() Transformation

WHERE +

GROUP BY \$_interval fill(null) routing-engine +

+ +

Visualization

Cancel Save

This example shows separate configurations for two devices – vMX1 and MX240 – in a single canvas titled **Canvas1**. The query section takes routing engine’s CPU utilization (re-cpu-utilization) as field input, computes the mean every minute for the three devices. These example configurations are set for the time series graph captured over 6 hours.

- In the **Query** section, the **FROM** fields shows the input for device-group name, device name, and rule, respectively.
- The **SELECT** fields contain inputs for a KPI key based on the rule that you set earlier. If the rule you selected is system.cpu/check-system-cpu, the **SELECT** field could show options for CPU utilization value, CPU utilization high threshold, CPU utilization low threshold, and so on.
- The **Aggregation** field describes data aggregation computation to be performed on the KPI field data (CPU utilization in this example) such as mean, median, sum, and so on.
- The **Transformation** field (optional) describes computations performed on subsequent values in a single field. For example, rate of change, elapsed time, differences, and so on.
- The **WHERE** section allows you to set a condition to filter the field data. You must enter a left operand, a right operand and a conditional operand such as **equal to**, **matches with**, **greater than**, **less than**, and so on. The left operand can be either a key or a field name.

In this example, no conditions are set in the **WHERE** section.

- The **GROUP BY** section allows you to group the filtered data by a time interval you specify and handle output value in cases where no field data is available. You can choose group by time interval from the drop-down list.

If you set the time interval in **GROUP BY** field as **\$_interval**, then Grafana sets the time interval for data aggregation based on the number of points received in TSDB.

- In the **Visualization** section, you can enter the y-axis minimum and maximum values, y-axis label and choose unit such as bytes, kilobytes, and so on.

In these example configurations, the **Visualization** section is not configured.

3. Sample graph panel output for two configurations:



The canvas titled **Canvas1** has two graphs (because of two separate configurations) for the vMX1 and MX240 devices side-by-side.

Grafana

IN THIS CHAPTER

- [Grafana Overview | 423](#)
- [Monitor Network Device Health Using Grafana | 424](#)

Grafana Overview

IN THIS SECTION

- [Access the Grafana UI and View Dashboards | 423](#)

Grafana is an open-source data visualization tool. You use the Grafana UI to create and to view charts, graphs, and other visuals to help organize and understand data. You can create dashboards to monitor the status of devices, and you can also query data and view the results from the UI. Grafana UI renders data from Paragon Automation time series database (TSDB). For more information, see [Grafana Documentation](#).

Access the Grafana UI and View Dashboards

You access the Grafana UI directly from the Paragon Automation UI.

All logged in users can access and use Grafana and its features by clicking **Monitoring > Graphs > Grafana** from the Paragon Automation UI.

You use dashboards to monitor the status and health of devices. Follow these steps to view existing dashboards or to create new dashboards from the Grafana UI:

- To create a dashboard, click **(+) > Create > Dashboards > Add Panel** from the left-nav panel.
- To view existing dashboards, click **Dashboards > Home** from the left-nav panel.

RELATED DOCUMENTATION

[Monitor Network Device Health Using Grafana | 424](#)

Monitor Network Device Health Using Grafana

SUMMARY

Read these topics to understand how to view pre-populated graphs and run queries from the Grafana UI.

IN THIS SECTION

- [Run a Query | 424](#)
- [View Pre-populated Graphs | 426](#)

You use Grafana to create and to view charts, graphs, and other visuals to help organize and understand data. For more information on Grafana and how to access the Grafana UI, see "[Grafana Overview](#)" on [page 423](#)

Run a Query

You can run queries from the Grafana UI. For more information, see [Query a Data Source](#).

When you create a chart, the data is rendered from the TSDB database. The TSDB database library (shim) is based on InfluxDB, an open-source time series database. Information sent to and received from the TSDB database is in the InfluxDB format. Paragon Automation also supports Grafana's open source InfluxDB plug-in. After you run a query from the Grafana UI, the open-source InfluxDB sends a back-end request to the Paragon Automation TSDB. The Paragon Automation TSDB sends the response in InfluxDB format.

In earlier releases, when you run a query from the Grafana UI using influxDB as a data source, you can only view data of a specific device in a device group. You use syntax parameters such as FROM and SELECT that the influxDB recognizes. You cannot view aggregate data of a device group, or view aggregate data of multiple device groups using the influxDB plug-in. However, Paragon Automation also supports the Juniper Paragon Insights TSDB plug-in.

This plug-in is installed when you install Paragon Automation. The Juniper Paragon Insights TSDB plug-in is based on the influxDB plug-in but also supports the ability to view aggregate data of a device group, or view aggregate data of multiple device groups. You can view aggregate data by using syntax parameters such as DEVICE, DEVICE GROUPS, and MEASUREMENT. The default data source for Paragon Insights in Grafana is based on the Juniper Paragon Insights TSDB plug-in.

For more information on supported syntax parameters, see [Table 76 on page 425](#).

Figure 28: Querying from Grafana UI




Table 76: Examples of InfluxDB and TSDB Syntax Parameters

InfluxDB Querying Options	TSDB Querying Options
FROM—Select the data source (device, device group, topic/rule).	DEVICE—Select one or more than one device for this query.
SELECT—Select the data field, and apply aggregation and transformation types to the data.	DEVICE GROUP—Select one or more than one device group for this query.
GROUP BY—Specify how to group data based on KPI keys.	MEASUREMENT—Select the Paragon Insights topic or rule name.
FORMAT AS—Specify if you want to format as table, time series, or log.	FIELDS—Select the fields to which you want to apply the aggregation.
	WHERE—Filter data based on tags and fields.
	GROUP BY—Specify how to group data based on KPI keys.
	FORMAT AS—Specify if you want to format as table, time series, or log.

Table 76: Examples of InfluxDB and TSDB Syntax Parameters (Continued)

InfluxDB Querying Options	TSDB Querying Options
	PER_DEVICE_LIMIT—Set the per-device limit of the base query for each device (database) that you have selected.
	AGGREGATED_LIMIT—Set the aggregate limit for the device (databases) after the base query is run.
	AGGREGATED_LIMIT—Rename a field. For example, if you select mean as an aggregation function for a field named bps , you can give it an alias name called mean_bps .

View Pre-populated Graphs

You can view pre-populated graphs from the Grafana dashboard.

You can view pre-populated graphs from the Dashboards section of the Grafana Home page. You can also view these pre-populated graphs by clicking **Dashboard>Manage>Paragon Insights Cluster Health**. The following graphs are available by default:

- CPU Usage—View cluster-wise or node-wise information on CPU (%) usage at different levels.
- Disk Read Usage (r_await)—View cluster-wise or node-wise information on average time taken for disk reads to be served.
- Disk Write Usage (w_await)—View cluster-wise or node-wise information on average time taken for disk writes to be served.
- Node Memory Available—View information on free memory available on a node in a selected cluster.

Click the name of a graph to view more information of that graph. You can view a pre-populated graph of all nodes in a cluster at a time, and also view pre-populated graphs of one or more nodes in a cluster.

7

PART

Manage Sensor Settings, Insights Settings, and Data Summarization Profiles

[Sensor Settings](#) | 428

[Insights Settings](#) | 613

[Data Summarization Profiles](#) | 647

Sensor Settings

IN THIS CHAPTER

- [Sensors Overview | 429](#)
- [About the Ingest Settings Page | 454](#)
- [Configure NetFlow Settings | 455](#)
- [Configure a Rule Using Flow Sensor | 461](#)
- [About the Frequency Profiles | 468](#)
- [Manage Frequency Profiles | 469](#)
- [Apply a Frequency Profile | 473](#)
- [Configure Offset Time | 474](#)
- [Configure a Rule Using Server Monitoring Sensor | 482](#)
- [Configure Native GPB Ingest | 485](#)
- [Configure sFlow Settings | 486](#)
- [Configure SNMP Ingest | 499](#)
- [Configure a Rule Using SNMP Scalar | 503](#)
- [Configure SNMP Trap and Inform Notifications | 504](#)
- [Configure Outbound SSH Port for iAgent | 515](#)
- [Configure System Log Ingest | 516](#)
- [System Log Optional Configurations | 525](#)
- [Configure a Rule Using Syslog | 526](#)
- [Understand Inband Flow Analyzer 2.0 | 532](#)
- [Configure Device Details for Inband Flow Analyzer Devices | 538](#)
- [Delete an Inband Flow Analyzer Device | 539](#)
- [Understand Bring Your Own Ingest | 540](#)
- [Load BYOI Default Plug-ins | 541](#)
- [Configure Bring Your Own Ingest Default Plug-in Instances | 542](#)
- [Build and Load BYOI Custom Plug-in Images | 544](#)
- [Configure Bring Your Own Ingest Custom Plug-in Instances | 555](#)

- [Use Sample Rule and Playbook Configurations for BYOI Custom Plug-in Instances | 557](#)
- [Configure Ingest Mapping for Default BYOI Plug-in Instances | 558](#)
- [Delete a BYOI Plug-in | 560](#)
- [About the Diagnostics Page | 561](#)
- [Use the Self Test Tool | 564](#)
- [Use the Reachability Test | 566](#)
- [Use the Ingest Test Tool | 567](#)
- [Use the No-Data Tool | 568](#)
- [Paragon Insights Tagging Overview | 570](#)
- [Types of Tagging | 577](#)
- [Add a Tagging Profile | 585](#)
- [Apply a Tagging Profile | 590](#)
- [Delete a Tagging Profile | 591](#)
- [Understand User-Defined Actions and Functions | 593](#)
- [Modify User-Defined Action, Function, and Workflow Engines | 594](#)
- [Enable UDA Scheduler in Trigger Action | 599](#)
- [Understand kube-state-metrics Service | 600](#)

Sensors Overview

IN THIS SECTION

- [Native GPB | 431](#)
- [NetFlow | 431](#)
- [sFlow | 433](#)
- [OpenConfig | 434](#)
- [Syslog | 438](#)
- [Server Monitoring Sensor | 440](#)
- [iAgent \(CLI/NETCONF\) | 446](#)
- [SNMP | 453](#)

Paragon Insights accepts data from Juniper, third-party devices, and from various types of telemetry sensors including traditional network management protocols like system log and SNMP. Paragon Insights supports push and pull models of data collection. In the push model, your devices push telemetry data to Paragon Insights through trap notifications, for instance. In the pull mode, Paragon Insights periodically polls your devices for data. This guide describes each of the supported ingest methods, with examples, sorted by whether they fall into the push or pull model. Along with each description, we provide the required Junos OS version and device configurations needed to enable the specific ingest type.

As the number of objects in the network, and the metrics they generate, have grown, gathering operational statistics for monitoring the health of a network has become an ever-increasing challenge. Traditional 'pull' data-gathering models, like SNMP and the CLI, require additional processing to periodically poll the network element, and can directly limit scaling.

The 'push' model overcomes these limits by delivering data asynchronously, which eliminates polling. With this model, the Paragon Insights server can make a single request to a network device to stream periodic updates. As a result, the 'push' model is highly scalable and can support the monitoring of thousands of objects in a network. Junos devices support this model in the form of the [Junos Telemetry Interface \(JTI\)](#).

Paragon Insights currently supports five push-model sensors:

- ["Native GPB" on page 431](#)
- ["NetFlow" on page 431](#)
- ["sFlow" on page 433](#)
- ["OpenConfig" on page 434](#)
- ["Syslog" on page 438](#)
- ["Server Monitoring" on page 440](#)
- ["Outbound SSH \(Device-Initiated\)" on page 451](#)

While the 'push' model is the preferred approach for its efficiency and scalability, there are still cases where the 'pull' data collection model is appropriate. Two examples might be when a device doesn't support the Junos Telemetry Interface (JTI), or when managing third party devices. With the pull model, Paragon Insights requests data from network devices at periodic, user-defined intervals.


Paragon Insights currently supports the following pull-model sensors:

- ["iAgent \(CLI/NETCONF\)" on page 446](#)
- ["SNMP" on page 453](#)

Native GPB

Native sensors use a Juniper-proprietary data model using Google Protocol Buffers (GPB). The device pushes telemetry data (when configured) over UDP.

The device pushes data from the Packet Forwarding Engine, that is, directly from a line card. This means telemetry data is sent over the forwarding plane, so the collector must have in-band connectivity to the device.




To use native format, you configure the device with settings that include where to send the telemetry data. When you configure Paragon Insights to start collecting the data, the stream is already flowing towards the server.

For more information on native sensors, see [Understanding the Junos Telemetry Interface Export Format of Collected Data](#).

NetFlow

Paragon Insights supports NetFlow v9 and NetFlow v10 (IPFIX) natively as NetFlow ingest method, using a data model that aligns with other Paragon Insights ingest mechanisms. NetFlow is a network protocol for collecting IP traffic statistics, which can then be exported to a tool for analysis. The NetFlow v9 data export format is described in [RFC 3954](#); NetFlow v10 is officially known as IPFIX and standardized in [RFC 7011](#).

Junos devices support flow monitoring and aggregation using these protocols; the Junos OS samples the traffic, builds a flow table, and sends the details of the flow table over a configured UDP port to a collector, in this case Paragon Insights. Paragon Insights receives the incoming Netflow data, auto-detects it as v9 or v10, and process it further.



As shown above, the network device pushes data from the Packet Forwarding Engine, that is, directly from a line card. This means flow data is sent over the forwarding plane, so the collector must have in-band connectivity to the device. To use the flow sensor option, you configure the device with settings that include where to send the flow data. When you configure Paragon Insights to start collecting the data, the flow data is already flowing towards the server.

Paragon Insights uses flow templates as a mechanism to identify and decode incoming flow data before sending it for further processing. Paragon Insights provides predefined flow templates for NetFlow v9 and v10 (IPFIX), or you can define your own. The predefined templates match those which the Junos OS currently supports. For example, the Junos OS template, `ipv4-template`, aligns with the Paragon Insights template `hb-ipfix-ipv4-template`. To view the fields used in the Junos OS templates, see [Understanding Inline Active Flow Monitoring](#).

NOTE: In the current ingest implementation for NetFlow, the following field types are not supported:

- Fields for enterprise specific elements
- Variable length fields



WARNING: For NetFlow ingest, ensure that there is no source NAT in the network path(s) between the device and Paragon Insights. If the network path contains source NAT, then the received device information is not accurate.

A typical workflow includes adding a NetFlow configured device in Paragon Insights, configuring NetFlow templates, configuring a rule with Flow sensor, and deploying a playbook with the rule to a device group.

- Configure devices to use NetFlow in Paragon Insights. See ["Edit Devices" on page 168](#).
- Add the device to a device group. See ["Add a Device Group" on page 177](#).
- Define NetFlow ingest settings. See ["Configure Netflow Settings" on page 455](#).
 - Use pre-defined NetFlow templates or
 - Create your own NetFlow template
 - Clone an existing NetFlow template
- Configure a rule that uses a flow sensor. See ["Configure a Rule Using Flow Sensor" on page 461](#).
- Add the rule to a playbook. ["Create a New Playbook Using the Paragon Insights GUI" on page 309](#)

- Deploy the playbook. ["Manage Playbook Instances" on page 312.](#)
- Monitor the device for NetFlow traffic

With the playbook applied, you can begin to monitor the devices.

1. Click **Monitoring > Network Health** in the left navigation bar and click on the **Device Group** tab.
2. Select the device group to which you applied the playbook from the **Device Group** pull-down menu.
3. Select one or more of the devices to monitor.
4. In the Tile View, the external tile contains the parameters from the rule you configured earlier.

sFlow

Paragon Insights supports sFlow (v5) natively as another flow-based ingest method.

sFlow is a statistical sampling-based technology for high-speed switched or routed networks. You can configure sFlow to continuously monitor traffic at wire speed on all interfaces simultaneously if you want.

sFlow provides or helps with:

- Detailed and quantitative traffic measurements at gigabit speeds
- Insight into forwarding decisions
- Troubleshooting for network issues
- Congestion control
- Security and audit trail analysis
- Route profiling

Everything that sFlow does above, it does without impact to forwarding or network performance. For more information on sFlow, see: [RFC 3176, InMon Corporation's sFlow: A Method for Monitoring Traffic in Switched and Routed Networks.](#)

As a statistical sampling protocol, Juniper's sFlow agent samples the traffic and counters on network interfaces, creates sFlow datagrams, and forwards them to external sFlow collectors. Paragon Insights is one such collector.

To know how to configure sFlow packets in Paragon Insights, go to ["Configure sFlow Settings" on page 486.](#)

OpenConfig

IN THIS SECTION

- [gNMI-Encoded OpenConfig RPC | 434](#)
- [Device Configuration for OpenConfig | 435](#)

To use OpenConfig format, you configure the device as a gRPC server. With Paragon Insights acting as the client, you define which sensors you want it to subscribe to, and it makes subscription requests towards the device.

Data streamed through gRPC is formatted in OpenConfig key/value pairs in protocol buffer (GPB) encoded messages. Keys are strings that correspond to the path of the system resources in the OpenConfig schema for the device being monitored; values correspond to integers or strings that identify the operational state of the system resource, such as interface counters. OpenConfig RPC messages can be transferred in bulk, such as providing multiple interface counters in one message, thereby increasing efficiency of the message transfer.

For more information on OpenConfig sensors, see [Understanding OpenConfig and gRPC on Junos Telemetry Interface](#).

gNMI-Encoded OpenConfig RPC

gNMI-encoded OpenConfig works much like OpenConfig RPC in that you must set the network device up as an OpenConfig server to which Paragon Insights makes subscription requests. However, gNMI supports more subscription types than Paragon Insights currently supports. Currently, Paragon Insights only supports gNMI STREAM subscriptions in the SAMPLE mode. STREAM subscriptions are long-lived subscriptions that continue, indefinitely, to transmit updates relating to the set of paths configured within the subscription. SAMPLE mode STREAM subscriptions must include a `sample_interval`.

The messages returned to the client through gNMI are encoded by the device in protobuf, JSON, or JSON-IETF format and cannot be sent in bulk. This, in part, makes gNMI-encoded messaging less efficient than gRPC-encoded messaging.



WARNING:

- For JSON or JSON-IETF, it is assumed that the device returns gNMI updates as only leaf values encoded in JSON, rather than returning an entire hierarchy or sub-hierarchy as a JSON object.
- Numbers encoded in JSON or JSON-IETF are decoded by Paragon Insights as either float64, int64, or string, according to [RFC 7159](#) and [RFC 7951](#). If your OpenConfig rules contain fields that are of a different type, we recommend that you change the field types accordingly.

Junos OS and Cisco devices can be managed by Paragon Automation using gNMI-encoded OpenConfig. If a device does not support gNMI in general, or the STREAM subscription in SAMPLE mode, or does not support an OpenConfig request, it returns one of the following errors:

- Unimplemented
- Unavailable
- InvalidArgument

In the case of a Junos OS or Cisco device, the error causes the connection to fall back to OpenConfig RPC. In the case of a third-party device, the connection simply fails due to the error.

gNMI-encoded OpenConfig can be enabled at the device or device-group level. If enabled at the device-group level, then all devices added to the group use gNMI by default. If enabled (or not enabled) at the device level, then the device level setting takes precedence over the device-group level setting.

NOTE: During the initial connection gNMI devices attempt to perform an initial sync with the client. The device sends a continuous stream of data until the device and the collector (Paragon Insights) are in sync. After initial sync, the device begins normal streaming operations based on the configured reporting rate. Because of the processing load this creates, Paragon Insights has this feature disabled by default. It can be enabled at the device-group or device level if needed.

For more information about gNMI, see: [gRPC Network Management Interface \(gNMI\)](#).

Device Configuration for OpenConfig

OpenConfig requires:

- Junos OS Version: 16.1 or later
 - The OpenConfig sensor requires that the Junos device have the OpenConfig and network agent packages installed. These packages are built into Junos OS Releases 18.2X75, 18.3, and later. For

releases between 16.1 and 18.2X75 or 18.2, you must install the OpenConfig and Network Agent packages.

Before you install the Network Agent package:

- Install Junos OS Release 16.1R3 or later.
- Install the OpenConfig for Junos OS module. Using a Web browser, navigate to the All Junos Platforms software download URL on the Juniper Networks webpage: <https://www.juniper.net/support/downloads/>. From the **Network Management** tab, scroll down to select **OpenConfig**. Select the **Software** tab. Select the **OpenConfig Package (Junos with upgraded FreeBSD)**. For more information, see *Installing the OpenConfig Package*.
- Install Secure Sockets Layer (SSL) certificates of authentication on your Juniper Networks device.

NOTE: Only server-based SSL authentication is supported. Client-based authentication is not supported.

An example of a valid Network Agent package name is:

network-agent-x86-32-16.1R4.12-C1.1.tgz

NOTE: Each version of the Network Agent package is supported on a single release of Junos OS only. The Junos OS version supported is identified by the Junos OS release number included in the Network Agent package name.

Use the 32-bit Network Agent package for both 32-bit and 64-bit versions of Junos OS or Junos OS Evolved.

To download and install the Network Agent package:

1. Navigate to the All Junos Platforms software download URL on the Juniper Networks webpage: <https://www.juniper.net/support/downloads/>.
2. Select the name of the Junos OS platform for the software that you want to download.
3. Select the release number (the number of the software version that you want to download) from the **Release** drop-down list to the right of the Download Software page.
4. Select the **Software** tab.
5. Navigate to the **Tools** section of the **Software** tab and select the **Junos Network Agent** package for the release.

6. Log in to the Juniper Networks authentication system using the username (generally your e-mail address) and password supplied by a Juniper Networks representative.
7. Download the software to a local host.
8. Copy the software to Juniper Networks device or to your internal software distribution site.
9. Install the new network-agent package on the device by issuing the `request system software add package-name` from the operational mode:

For example:

```
user@host > request system software add network-agent-x86-32-16.1R3.16-C1.0.tgz
```

NOTE: The command uses the `validate` option by default. This option validates the software package against the current configuration as a prerequisite to adding the software package to ensure that the device reboots successfully. This is the default behavior when the software package being added is a different release.

To verify whether you have the OpenConfig and the Network Agent packages installed, enter the following command:

```
user@host> show version | match "Junos:|openconfig|na telemetry"

Junos: 19.2R1.8
JUNOS na telemetry [19.2R1.8]
JUNOS Openconfig [19.2R1.8]
```

See [Understanding OpenConfig and gRPC on Junos Telemetry Interface](#) for more information.

- Network agent is not supported on PPC platforms (MX104, MX80, and so on).

Device Configuration

Configure a device by entering the following command:

```
set system services extension-service request-response grpc clear-text port <port number>
```

To configure OpenConfig port under device configuration in Paragon Automation GUI, see Editable Fields on the Edit Devices Page table in ["Edit Devices" on page 168](#).


Syslog

IN THIS SECTION

- [Requirements to configure Syslog in Paragon Insights | 439](#)

In addition to the JTI-related options above, Paragon Insights also supports system logs as a data collection method, using a data model that aligns with other ingest mechanisms.

A device can push syslog messages (when configured) over UDP to the Paragon Insights server either out-of-band through the Routing Engine (RE) using the router's management interface, or in-band through the Packet Forwarding Engine, that is, directly from a line card.



To use syslog format, you configure the device with settings that include where to send the syslog messages. When you configure Paragon Insights to start collecting the data, messages are already flowing towards the server.

For more information on syslog as used by Juniper devices, see [Junos OS System Log Overview](#).

A syslog message consists of a header, structured data in key-value format within square brackets, and the log message. The header consists of the following information:

- Log priority
- Version number of Syslog protocol specification on header format.
Currently, this number is 1.
- Timestamp of when the message was generated in the format of "Mmm dd hh:mm:ss".
- Hostname identifies the device that sent the syslog message.
- Application name
- Application process ID

- Message ID

Requirements to configure Syslog in Paragon Insights

Syslog ingest requires some setup before you can use it as a sensor in a rule:

- Pattern - A pattern identifies some syslog event; you create a pattern for each event you want to monitor. You can configure patterns for both structured and unstructured events.
- Pattern set - With the patterns configured, you then group them into a pattern set, which you then reference when defining the syslog sensor settings within a rule.

Before you configure Patterns and Pattern Set for Syslog ingest, note that the following fields are common in syslog messages. Paragon Insights extracts these fields and includes them automatically in the raw table, enabling you to make use of them directly when creating a rule, and avoiding the need to configure patterns.

To illustrate use of these values, consider the following example syslog messages:

Structured - <30>1 2019-11-22T03:17:53.605-08:00 R1 mib2d 28633 SNMP_TRAP_LINK_DOWN [junos@2636.10.1.1.2.29 snmp-interface-index="545" admin-status="up(1)" operational-status="down(2)" interface-name="ge-1/0/0.16"] ifIndex 545, ifAdminStatus up(1), ifOperStatus down(2), ifName ge-1/0/0.16

Equivalent unstructured - <30>Nov 22 03:17:53 R1 mib2d[28633]: SNMP_TRAP_LINK_DOWN: ifIndex 545, ifAdminStatus up(1), ifOperStatus down(2), ifName ge-1/0/0.16

System-generated fields:

- "__log_priority__" - Priority of syslog message
 - In the examples, <30> denotes the priority
- "__log_timestamp__" - Timestamp in epoch in the syslog message
 - In the structured example, 2019-11-22T03:17:53.605-08:00 is converted to epoch with **-08:00** indicating the time zone
 - In the unstructured example, the time zone from the configuration will be used to calculate epoch
- "__log_host__" - Host name in the syslog message
 - In the examples, R1 denotes the host name
- "__log_application_name__" - Application name in the syslog message
 - In the examples, mib2d is the application name
- "__log_application_process_id__" - Application process ID in the syslog message

- In the examples, 28633 is the ID
- `"__log_message_payload__"` - Payload in the message
 - Structured example - `"SNMP_TRAP_LINK_DOWN [junos@2636.10.1.1.2.29 snmp-interface-index="545" admin-status="up(1)" operational-status="down(2)" interface-name="ge-1/0/0.16"] ifIndex 545, ifAdminStatus up(1), ifOperStatus down(2), ifName ge-1/0/0.16"`
 - Unstructured example - `"SNMP_TRAP_LINK_DOWN: ifIndex 545, ifAdminStatus up(1), ifOperStatus down(2), ifName ge-1/0/0.16"`
- "Event-id" - Denotes the event ID configured in the pattern
 - In the examples, `SNMP_TRAP_LINK_DOWN` is the event ID

NOTE: Be sure not to define any new fields using a name already defined above.

To know how to configure Syslog ingest, see ["Configure System Log Ingest" on page 516](#).

Server Monitoring Sensor

In Paragon Automation, the Server Monitoring sensor collects data from servers and virtual machines on which you host the Paragon application. The sensor uses the third-party plug-in, Node Exporter. The Node Exporter plug-in is pre-installed in all server clusters of Paragon Automation. In the GUI, the default servers and virtual machines deployed in the Paragon Automation cluster are represented as devices that are automatically added to the **Paragon-Cluster** device group. The sensor collects data from servers and virtual machines to track CPU, memory, network, traffic, disk, and filesystem metrics. It writes the output to a time series database.

NOTE: Users must not delete the default **Paragon-Cluster** device group.

Paragon Automation has the following pre-configured playbooks to monitor server data.

- CPU utilization
- Disk reads and writes
- Errors, available bytes, and utilized bytes in filesystem
- Utilized bytes and available bytes in memory
- Received and transmitted total packet size, errors in received and transmitted packets, total received and transmitted multicast packets in network

When you configure a rule using Server Monitoring ingest, you can use the some of the sensor paths listed in [Table 77 on page 441](#).

Table 77: Server Metrics

Sensor Path	Description
<code>/node/boot/time/seconds</code>	Boot time in each server node.
<code>/node/cpu/seconds/total</code>	The total time (in seconds) the CPU stays in idle, system, user, and nice modes.
<code>/node/disk/read/bytes/total</code>	The total number of bytes read successfully.
<code>/node/disk/read/errors/total</code>	The total number of read errors in a node.
<code>/node/disk/read/retries/total</code>	The number of times the ingest tries to read from the disk if there is a failure.
<code>/node/disk/read/sectors/total</code>	The total number of sectors read successfully.
<code>/node/disk/read/time/seconds/total</code>	The total time taken to complete reads successfully per node.
<code>/node/disk/reads/completed/total</code>	The total number of reads completed successfully.
<code>/node/disk/write/errors/total</code>	The total number of errors in writes
<code>/node/disk/write/retries/total</code>	The number of times the ingest tries to write on the disk if there is a failure.
<code>/node/disk/write/time/seconds/total</code>	The total time taken to complete all writes.

Table 77: Server Metrics (Continued)

Sensor Path	Description
<code>/node/disk/writes/completed/total</code>	The total number of writes completed per node.
<code>/node/disk/written/bytes/total</code>	The total number of bytes written successfully.
<code>/node/disk/written/sectors/total</code>	The total number of sectors written successfully.
<code>/node/exporter/build/info</code>	A metric that has the value '1' and has version, revision, go version, and branch from which node exporter is built.
<code>/node/filesystem/avail/bytes</code>	The filesystem size available to non-root users.
<code>/node/filesystem/device/error</code>	The number of I/O errors that occur when collecting data from a filesystem.
<code>/node/filesystem/files</code>	The total number of index nodes permitted in a node.
<code>/node/filesystem/files/free</code>	The number of index nodes that are free for use in a node.
<code>/node/filesystem/free/bytes</code>	The free space (in bytes) available for the user, excluding reserved blocks.
<code>/node/filesystem/readonly</code>	Data that shows if the filesystem in a node is mounted as read-only.
<code>/node/filesystem/size/bytes</code>	The size of all files in bytes.

Table 77: Server Metrics (Continued)

Sensor Path	Description
/node/load1	Load on each server/host node captured every 1 minute.
/node/load15	Load on each server/host node captured every 15 minutes.
/node/load5	Load on each server/host node captured every 5 minutes.
/node/memory/active/bytes	Memory bytes that are actively used by processes.
/node/memory/compressed/bytes	Total size of compressed memory.
/node/memory/free/bytes	Total memory in bytes that is free for use in a node.
/node/memory/inactive/bytes	Memory bytes that are not actively used by processes.
/node/memory/swap/total/bytes	Total memory swapped in a node.
/node/memory/swap/used/bytes	The amount of swapped memory used in a node.
/node/memory/swapped/in/bytes/total	Total swapped in memory in a node.
/node/memory/swapped/out/bytes/total	Total swapped out memory in a node.
/node/memory/total/bytes	Total bytes of memory in a node.

Table 77: Server Metrics *(Continued)*

Sensor Path	Description
/node/memory/wired/bytes	Memory that cannot be swapped out.
/node/network/receive/bytes/total	Total size of packets received by a device.
/node/network/receive/errs/total	Total number of errors encountered by a device when receiving packets.
/node/network/receive/multicast/total	Total number of multicast packets received by a device.
/node/network/receive/packets/total	Total number of packets received by a device.
/node/network/transmit/bytes/total	Total size of packets sent from a device.
/node/network/transmit/errs/total	Total number of errors encountered a device when transmitting packets.
/node/network/transmit/multicast/total	Total number of multicast packets transmitted by a device.
/node/network/transmit/packets/total	Total number of packets transmitted by a device.
/node/scrape/collector/duration/seconds	Time taken by each collector to scrape metrics.
/node/scrape/collector/success	Number of times Node Exporter collector successfully scraped targets.

Table 77: Server Metrics (*Continued*)

Sensor Path	Description
<code>/node/textfile/scrape/error</code>	Errors encountered by Node Exporter when scraping targets using textfile scripts.
<code>/node/time/seconds</code>	Displays system time in seconds in the node since epoch (1970).
<code>/node/uname/info</code>	Name of the node from which Node Exporter collects metrics.

The following tags such as `mode`, `device` etc. can be used as key fields applicable to all metrics listed under main metrics (`/node/cpu` or `/node/network`).

When you configure a key field in a rule, you can mention only the key field name in the **Path** field.

- `/node/cpu/`
 - **cpu**: The number of cores available in CPU.
 - **mode**: The type of CPU usage in a node such as `idle`, `system`, `user`, and `nice`.
- `/node/disk/`
 - **device**: Name of disks such as `disk0`, `disk1`, `sda`, `sdb`, or `sd`.
- `/node/filesystem/`
 - **device**: Disk paths such as `/dev/sda1`, `/dev/sda2`, and `/dev/sdb1`
 - **fstype**: Type of partition formatting such as `ext4`, `NTFS` (New Technology File System), and `APFS` (Apple File System).
 - **mountpoint**: Mount paths in the device.
- `/node/network/`
 - **device**: Interface names of the device such as `wlan0`, `en0`, or `docker0`.

To configure an example rule using Server Monitoring ingest, see "[Configure a Rule Using Server Monitoring Sensor](#)" on page 482

iAgent (CLI/NETCONF)


IN THIS SECTION

- Define PyEZ Table/View | 447
- Gather Output from Device | 448
- Generate JSON for Use in Paragon Automation Database | 449
- Outbound SSH (Device-Initiated) | 451

For all the benefits of the 'push' data collection methods, some operational and state information is available only through CLI/VTY commands. iAgent fills this gap by taking advantage of NETCONF/SSH functionality to provide Paragon Automation with the ability to connect to a device, run commands, and capture the output.

iAgent sensors use NETCONF/SSH and YAML-based PyEZ tables and views to fetch the necessary data. Both structured (XML) and unstructured (VTY commands and CLI output) data is supported.

With iAgent, the Paragon Automation server initiates inbound SSH requests over any available network interface, whether in-band or out-of-band; and the device responds (when properly configured) with the requested data.



g-300770

In Paragon Automation Platform, the iAgent ingest is extended to other vendor devices from Arista Networks, Palo Alto Networks, and Cisco.

You must configure a Junos device to send telemetry data using iAgent sensor or a NETCONF connection.

At minimum, iAgent (NETCONF) requires:

- Junos OS Version: 11.4 or later

- Minimum required device configuration:

```
set system services netconf ssh
```

To configure iAgent or NETCONF port for inbound connection in Paragon Automation GUI, see [Table 46 on page 168](#).

Paragon Automation uses Netmiko to make inbound SSH connections over the selected port to a third-party device. To gather device information, Paragon Automation sends CLI commands over SSH and receives string blobs back as output. The string blobs are then parsed through TextFSM, using NTC templates into JSON format and then stored in the database. Default templates are located at `/srv/salt/_textfsm`. You can visit the GitHub repository of [NTC Templates](#) for network devices.

For advanced users who need a template which does not exist, you can create your own templates and upload them to Paragon Insights using the **Upload Rule Files** button on the **Configuration > Rules** page. User defined templates are stored at `/jfit/_textfsm`. The files must end with the `.textfsm` suffix. To know how to upload pre-defined rules in Paragon Automation Platform, see ["Add a Predefined Rule" on page 340](#).

TextFSM is integrated into PyEZ's table/view feature which is an integral part of iAgent.

Example: PaloAlto Panos– Show Running Security Policy

To see the running security policy on a Panos device, we need to:

- ["Define PyEZ Table/View" on page 447](#)—Define a table/view for it
- ["Gather Output from Device" on page 448](#)—Gather the output by sending the needed CLI command to the device over SSH
- ["Generate JSON for Use in Paragon Automation Database" on page 449](#)—Generate JSON to store in Paragon Automation database

Define PyEZ Table/View

We need to define a PyEZ table that is used by the iAgent rule assigned to the Panos device. The following table definition lacks a view definition. Because of this, the entire output from the `show running security-policy` ends up getting stored in the database after processing.

```
---
PanosSecurityPolicyTable:
  command: show running security-policy
  platform: paloalto_panos
```

```
key: NAME
use_textfsm: True
```

(Optional) To store only a portion of the received data in Paragon Automation, you can define a view in the same file. The view tells Paragon Automation which fields to pay attention to.

```
---
PanosSecurityPolicyTable:
  command: show running security-policy
  platform: paloalto_panos
  key: NAME
  use_textfsm: True
  view: TrafficAndActionView

TrafficAndActionView:
  fields:
    source: SOURCE
    destination: DESTINATION
    application_service: APPLICATION_SERVICE
    action: ACTION
```

Gather Output from Device

Using an iAgent rule that references the PyEZ table (or table/view) defined above, Paragon Automation sends the command `show running security-policy` to the device which produces the following output:

```
"intrazone-default; index: 1" {
  from any;
  source any;
  source-region none;
  to any;
  destination any;
  destination-region none;
  category any;
  application/service 0:any/any/any/any;
  action allow;
  icmp-unreachable: no
  terminal yes;
  type intrazone;
}
```

```

"interzone-default; index: 2" {
    from any;
    source any;
    source-region none;
    to any;
    destination any;
    destination-region none;
    category any;
    application/service 0:any/any/any/any;
    action deny;
    icmp-unreachable: no
    terminal yes;
    type interzone;
}

dynamic url: no

```

Generate JSON for Use in Paragon Automation Database

Since the device configuration specifies Palo Alto Networks as the vendor and Panos OS as the operating system, the TextFSM template used for this example would look like this:

```

Value Key,Filldown NAME (.*)
Value Required FROM (\S+)
Value SOURCE (\S+)
Value SOURCE_REGION (\S+)
Value TO (\S+)
Value DESTINATION ([\S+\s+]*)
Value DESTINATION_REGION (\S+)
Value USER (\S+)
Value CATEGORY (\S+)
Value APPLICATION_SERVICE ([\S+\s+]*)
Value ACTION (\S+)
Value ICMP_UNREACHABLE (\S+)
Value TERMINAL (\S+)
Value TYPE (\S+)

Start
  ^${NAME}\s+\{
  ^\s+from\s+${FROM};

```

```

^\s+source\s+${SOURCE};
^\s+source-region\s+${SOURCE_REGION};
^\s+to\s+${TO};
^\s+destination\s+${DESTINATION};
^\s+destination-region\s+${DESTINATION_REGION};
^\s+user\s+${USER};
^\s+category\s+${CATEGORY};
^\s+application/service\s+${APPLICATION_SERVICE};
^\s+action\s+${ACTION};
^\s+icmp-unreachable:\s+${ICMP_UNREACHABLE}
^\s+terminal\s+${TERMINAL};
^\s+type\s+${TYPE};
^} -> Record

```

When the template above is used by Paragon Automation to parse the output shown previously, the resulting JSON looks like:

```

{"interzone-default; index: 2": {'ACTION': 'deny',
                                'APPLICATION_SERVICE': '0:any/any/any/any',
                                'CATEGORY': 'any',
                                'DESTINATION': 'any',
                                'DESTINATION_REGION': 'none',
                                'FROM': 'any',
                                'ICMP_UNREACHABLE': 'no',
                                'SOURCE': 'any',
                                'SOURCE_REGION': 'none',
                                'TERMINAL': 'yes',
                                'TO': 'any',
                                'TYPE': 'interzone',
                                'USER': ''},
 "intrazone-default; index: 1": {'ACTION': 'allow',
                                  'APPLICATION_SERVICE': '0:any/any/any/any',
                                  'CATEGORY': 'any',
                                  'DESTINATION': 'any',
                                  'DESTINATION_REGION': 'none',
                                  'FROM': 'any',
                                  'ICMP_UNREACHABLE': 'no',
                                  'SOURCE': 'any',
                                  'SOURCE_REGION': 'none',
                                  'TERMINAL': 'yes',
                                  'TO': 'any',

```

```
'TYPE': 'intrazone',
'USER': ''}]}
```

Outbound SSH (Device-Initiated)

Paragon Automation also supports iAgent (NETCONF) connections that are device-initiated using outbound SSH. This configuration makes Paragon Automation act as the client to the device making the connection. This type of connection is useful in environments in which the remote devices cannot accept incoming connections. All existing iAgent rules can be used when outbound SSH is configured in Junos devices.

NOTE: NETCONF over SSH connections are supported only in Junos devices.

Outbound SSH is disabled by default. You can configure an outbound SSH connection:

- In the ingest by configuring a single port. This port is used by all device groups.
- In device-groups by configuring its ports. This configuration takes precedence over the ingest configuration.

When you configure outbound SSH in device-groups, you must enter a TCP port number over which all the member devices initiate their NETCONF connections to Paragon Automation. You can disable outbound SSH in a device through management CLI. To configure outbound SSH ports in device groups, see ports section of the device group configuration in ["Add a Device Group" on page 177](#).

Paragon Automation supports a single TCP port for iAgent (NETCONF) outbound SSH connections from all device groups. This port can be configured at the ingest level. You can avoid opening multiple TCP ports in different device groups and simplify network management with a single port. To configure iAgent port at the ingest, see ["Configure Outbound SSH Port for iAgent" on page 515](#).

You can connect a device that is managed in different device groups through outbound SSH by configuring multiple clients, where each client has the same port. In this case, you must create as many copies of the device as there are device groups. Each device must have the same port number.

As an example, consider device r0 (10.1.1.1) configured for device groups *dg1* and *dg2*. To connect 10.1.1.1 to both device groups via the same outbound SSH port, you can create one more device r1 (10.1.1.1) with the same IP and deploy it in dg2.

You must also configure Paragon Automation for these ports in the respective device-groups. [Figure 29 on page 452](#) is an example device group configuration.

Figure 29: Edit Device Group Configuration

EDIT DEVICE GROUP ✕

Description ⓘ

Devices* ⓘ

▼ **Advanced**

Flow Ingest Deploy Nodes ⓘ

Reports ⓘ

Ingest Frequency Profiles ⓘ

Retention Policy ⓘ

> **Logging**

> **Notifications**

▼ **Ports**

Native Ports ⓘ ✓

Flow Ports ⓘ

sFlow Ports ⓘ

Syslog Ports ⓘ

SNMP Notification Ports ⓘ

Outbound SSH Ports ⓘ ✓

> **Summarization**

> **Data Rollup Summarization**

Cancel **Save** **Save and Deploy**

Using the following sample client configurations, device 10.1.1.1 can connect to two device groups using two outbound SSH clients with the same port.

```
set system services outbound-ssh client outbound-ssh1 device-id r0

set system services outbound-ssh client outbound-ssh1 10.1.1.1 port 2020

set system services outbound-ssh client outbound-ssh2 device-id r1
```

```
set system services outbound-ssh client outbound-ssh2 10.1.1.1 port 2020
```

NOTE: The 10.1.1.1 in the example denotes Paragon Automation (host) IP address.


SNMP

IN THIS SECTION

- [What's Next | 0](#)

SNMP is a widely known and accepted network management protocol that many network device manufacturers, including Juniper Networks, provide for use with their devices. It is a polling type protocol where network devices that are configured to use SNMP make configuration, diagnostic, and event information available to collectors, which must also be properly configured and authenticated. The collectors poll devices by sending specifically structured requests, called get requests, to retrieve data.

Paragon Automation supports SNMP as a sensor type, using standard get requests to gather statistics from the device. Paragon Automation makes requests over any available interface, whether in-band or out-of-band, and the device responds (when configured) with the requested data.



For information about SNMP as used on Junos OS devices, see [Understanding SNMP Implementation in Junos OS](#).

Paragon Insights also supports scalar object instances along with tabular objects in SNMP.

- The SNMP object can be scalar, tabular, or a combination of both in rules. When you create a rule using SNMP ingest, you can add:
 - Only scalar fields.
 - A combination of tabular and scalar fields.

- A tabular column along with the index queried as a scalar object.

A tabular column queried as a scalar comes with the limitation that the index number does not refer to the same Object across all the devices when you configure the tabular field in rule. For example, **IF-MIB::ifAdminStatus.16**. The ifAdminStatus is a column in **IF MIB** table. The **IF-MIB::ifAdminStatus.16** refers to the table column with index 16.

- Only tabular fields.
- A scalar object is identified by its MIB name (for example, **JUNIPER-MIB::scalarObjectName**) or as an OID.
- Paragon Insights validates a given scalar by checking the *MAX-ACCESS* property in the MIB definition.

If you find *MAX-ACCESS* in the MIB definition set to read-only, read-create, or read-write, then that object can be queried as a scalar.

The complete path to query a scalar object is **MIB-Name::table column name:index number**.

For example, **IF-MIB::ifInOctets.16**.

WHAT'S NEXT

["Configure a Rule Using SNMP Scalar" | 503](#)

["Configure SNMP Ingest" | 499](#)

["Configure SNMP Trap and Inform Notifications" | 504](#)

About the Ingest Settings Page

IN THIS SECTION

- [Tasks You can Perform | 455](#)

You can use the Ingest Settings page to define NetFlow ingest settings, implement log analysis using log patterns and pattern sets within the syslog ingest settings, apply tags to data at the ingest level before Paragon Insights processes the telemetry data, set frequency profiles to manage sensor and time

frequencies in rules in a centralized method, set port for Native GPB sensor, and configure sFlow settings.

To access the Ingest Settings page in Paragon Automation GUI, go to **Configuration > Data Ingest > Settings**.

Tasks You can Perform

The following are the tasks you can perform in the Ingest Settings page:

- Configure NetFlow Template Settings. See "[Configure NetFlow Settings](#)" on page 455.
- Configure Syslog for log analysis. See "[Configure System Log Ingest](#)" on page 516.
- Add frequency profiles. See "[Manage Frequency Profiles](#)" on page 469.
- Add tagging profiles. See "[Add a Tagging Profile](#)" on page 585.
- Configure Native GPB port. See "[Configure Native GPB Ingest](#)" on page 485.
- Configure sFlow ingest. See "[Configure sFlow Settings](#)" on page 486.
- Configure SNMP. See "[Configure SNMP Ingest](#)" on page 499.
- Configure SNMP trap and inform notifications. See "[Configure SNMP Trap and Inform Notifications](#)" on page 504
- Configure outbound SSH TCP port. See "[Configure Outbound SSH Port for iAgent](#)" on page 515.
- Configure ingest mappings for default Bring Your Own Ingest plug-ins. See "[Configure Ingest Mapping for Default BYOI Plug-in Instances](#)" on page 558.

RELATED DOCUMENTATION

| [About the Diagnostics Page](#) | 561

Configure NetFlow Settings

IN THIS SECTION

● [Use Pre-defined NetFlow Templates](#) | 456

- [Create Custom NetFlow Templates | 456](#)
- [Delete a NetFlow Template | 457](#)
- [Clone an Existing NetFlow Template | 458](#)
- [Configure Flow Source IP Address | 459](#)
- [Configure Flow Ports | 460](#)

Use Pre-defined NetFlow Templates

NetFlow templates provide a mechanism to identify and decode incoming flow data before sending it for further processing within Paragon Insights.

1. Click **Configuration > Data Ingest > Settings** in the left-nav bar.
2. Click the **NetFlow** tab on Ingest Settings page.
3. On the NetFlow settings page, review the available templates for use in a rule.

Usage Notes:

- Notice that there are default flow templates for IPv4, IPv6, MPLS, MPLS-IPv4, MPLS-IPv6, and VPLS, for each of NetFlow v9 and v10.
- The NetFlow templates include recognition patterns, called include fields and exclude fields, which help to recognize, identify, and categorize the incoming messages.
- Since NetFlow messages don't distinguish between keys and values (all fields are simply incoming data), the templates specify which fields should be treated as keys for raw data.

Create Custom NetFlow Templates

If the existing templates do not meet your needs, you can create your own template. You can also use custom templates to support other vendors' devices.

1. On the Netflow settings page, click the plus (+) icon.
2. In the Add Template window that appears, fill in the following fields (you can leave the other settings as is):
 - Template Name—Give the template a name
 - Description—Provide a description for the template
 - NetFlow version—Select **Netflow v9** or **Netflow v10**
 - Priority—Select any value from 1 through 10 depending on the level of priority

- Include Fields—Add one or more fields that you want included in the template you wish to use
- Exclude Fields—Add one or more fields that you do not want included in the template you wish to use
- Key Fields—Specify which fields in the incoming messages should be treated as keys

3. Click **Save & Deploy**

You should now see the template added to the NetFlow settings page.

4. (Optional) Repeat the steps above to create more templates.

Usage Notes:


- Priority - when a playbook includes multiple rules using the flow sensor, the priority value identifies which sensor and template gets priority over the other(s).
- Include/Exclude fields - include fields to help identify the template to use, or at least a 'short list' of templates to use; exclude fields then narrow down to the single desired template.
 - Example 1 - consider the *hb-ipfix-ipv4-template* template: it includes two IPv4 fields to narrow down to *hb-ipfix-ipv4-template* and *hb-ipfix-mpls-ipv4-template*, and excludes an MPLS field to eliminate *hb-ipfix-mpls-ipv4-template*, leaving only *hb-ipfix-ipv4-template*.
 - Example 2 - consider the *hb-ipfix-mpls-ipv4-template* template: it includes the same two IPv4 fields to narrow down to *hb-ipfix-ipv4-template* and *hb-ipfix-mpls-ipv4-template*. It also includes an MPLS field, which immediately eliminates the former template and leaving the latter as the template to use.

Delete a NetFlow Template

To delete a NetFlow template:

1. Click **Configuration > Data Ingest > Settings** from the left-nav bar.
The Ingest Settings page is displayed.
2. Click the **NetFlow** tab to view the NetFlow Settings page.
3. Select the template that you want to delete, and click the **delete (trash can)** icon.
The **CONFIRM DELETE TEMPLATE** pop-up appears.
4. Do one of the following:

Figure 30: Confirm Delete Template Pop-up



- Click **Yes** to delete the template from the database. However, the changes are not applied to the ingest service.

NOTE:

- We recommended that you do not delete a NetFlow setting that is currently in use.
- After you delete a particular NetFlow setting from the database, you cannot configure that NetFlow setting in new devices or device groups even if you have not deployed changes.
- You can also deploy changes to the ingest service or roll back the changes that you have already deleted, from the **Health Configuration Deployment Status** page. For more information, see ["Commit or Roll Back Configuration Changes in Paragon Insights" on page 185](#).

- Select the **Deploy changes** check box and then click **Yes** to delete the template from the database, and to apply the changes to the ingest service.
- (Optional) Click **No** to cancel this operation.

The NetFlow template is deleted.

Clone an Existing NetFlow Template

To clone an existing NetFlow template:

1. Click **Configuration > Data Ingest > Settings** from the left-nav bar.

The Ingest Settings page is displayed.

2. Click the **NetFlow** tab to view the NetFlow Settings page.
3. To clone a particular template, click **Clone**.

The Clone Template: *<name of template>* page is displayed.

From the Clone Template: *<name of template>* page, you can

- Edit the **Name**, **Description**, and **Priority** sections.
 - Choose between **Netflow v9** or **Netflow v10** versions.
 - Add or exclude fields from **Include Fields**, **Exclude Fields**, and **Key Fields**.
4. After you have made the necessary edits, click **Save** to save the modifications and to clone the template.
Alternatively, click **Save & Deploy** to save modifications, clone the template, and deploy the template.

Configure Flow Source IP Address

The raw flow data that Paragon Insights receives is in binary format and unreadable. In order to make this data usable, Paragon Insights processes the incoming flow data as follows:

- Paragon Insights listens for incoming flow data on a configured port
- Since NetFlow messages don't include a field that identifies the sending device, Paragon Insights uses the configured source IP address to derive a device ID.
- Templates identify and decode incoming flow data to determine which fields it contains

The resulting decoded and normalized data is now in a readable and usable format.

- Paragon Insights then performs further tagging, normalization, and aggregation as defined in the corresponding rule by the user.
- Finally, the time-series database (TSDB) receives the data. This is where things like trigger evaluation happen.



WARNING: For NetFlow ingest, ensure that there is no source NAT in the network path(s) between the device and Paragon Insights. If the network path contains source NAT, then the received device information is not accurate.

To configure source IP addresses in Device configuration:

1. Go to **Configuration > Devices**.

You are taken to the Device page.

2. Select a device you want to configure to send Flow data and click the edit button (pencil icon).
You are taken to the Edit *Device-Name* page.
3. Click the **Device ID Details** caret and enter the source IP address(es) in the Flow Source IP field.
If you want to enter multiple source IP addresses, separate each one with a comma.
4. Click **OK**.

To configure source IP addresses in Device Group configuration:

1. Go to **Configuration > Device Groups**.
You are taken to the Device Group configuration page.
2. Select a device group you want to configure to send Flow data and click the edit button (pencil icon).
You are taken to the Edit Device Group page.
3. Click the **Advanced** caret and enter the source IP address(es) in the Flow Ingest Deploy Nodes field.
If you want to enter multiple source IP addresses, separate each one with a comma.
4. Click **Save & Deploy**.

Configure Flow Ports

To configure Flow ports in Device Groups:

1. Go to **Configuration > Device Groups**.
You are taken to the Device Group configuration page.
2. Select a device group you want to configure to send Flow data and click the edit button (pencil icon).
You are taken to the Edit Device Group page.
3. Click the **Advanced > Ports** caret and enter the NetFlow sensor receiver ports in the Flow Ports field.
If you want to enter multiple ports, separate each one with a comma.
4. Click **Save & Deploy**.

RELATED DOCUMENTATION

[Configure a Rule Using Flow Sensor | 461](#)

[Configure sFlow Settings | 486](#)

Configure a Rule Using Flow Sensor

With the flow ingest settings complete, you can now create a rule using flow as the sensor.

This example rule includes three elements:

- A flow sensor that uses the NetFlow v10 IPv4 template
- Six fields capturing data of interest
- A trigger that indicates when traffic flow is higher or lower than expected

NOTE: See the usage notes at the end of this section for more detail on what has been configured.

1. Click **Configuration > Rules** in the left-navigation bar.
2. On the Rules page, click the **+ Add Rule** button.
The Rules page refreshes to show a nearly empty rule on the right part of the page.
3. In the top row of the rule window, leave the topic set as *external* and set the rule name that appears after the slash (/). In this example, it is *periodic-aggregation-flow-rule*.
4. Add a description and synopsis if you wish.
5. Click the **+ Add Sensor** button and enter the following parameters in the Sensors tab:

The screenshot shows the 'Sensors' configuration interface. At the top, there are tabs for 'Sensors', 'Fields', 'Vectors', 'Variables', 'Functions', 'Triggers', and 'Rule Properties'. The 'Sensors' tab is selected. Below the tabs, there is a '+ ADD SENSOR' button and a 'DELETE IPV4-FLOW-SENSOR' button. A list of sensors is shown, with 'ipv4-flow-sensor' selected. The configuration form for this sensor includes the following fields:

- Sensor Name ***: A text input field containing 'ipv4-flow-sensor'.
- Sensor Type**: A dropdown menu with 'Flow' selected.
- Template Name ***: A dropdown menu with 'hb-ipfix-ipv4-template' selected.

6. Now move to the Fields tab, click the **+ Add Field** button and enter the following parameters to configure the first field, *source-ipv4-address*:

Sensors Fields Vectors Variables Functions Triggers Rule Properties

+ ADD FIELD

source-ipv4-address DELETE SOURCE-IPV4-ADDRESS

Field Name * ?
source-ipv4-address

Description ?
Source IPv4 address

Field Type
string

Add to Rule Key ?

Ingest type (Field source)
Sensor

Sensor Path * Data if missing
ipv4-flow-sensor sourceIPv4Address Zero suppression Default value

Where (filter using expression)
+ ADD EXPRESSION

7. Click the **+ Add Field** button again and enter the following parameters to configure the second field, *destination-ipv4-address*:

Sensors Fields Vectors Variables Functions Triggers Rule Properties

+ ADD FIELD

source-ipv4-address DELETE DESTINATION-IPV4-ADDRESS

destination-ipv4-address

Field Name * ?
destination-ipv4-address

Description ?
Destination IPv4 Address

Field Type
string

Add to Rule Key ?

Ingest type (Field source)
Sensor

Sensor Path * Data if missing
ipv4-flow-sensor destinationIPv4Address Zero suppression Default value

Where (filter using expression)
+ ADD EXPRESSION

8. Click the **+ Add Field** button again and enter the following parameters to configure the third field, *sensor-traffic-count*:

Sensors **Fields** Vectors Variables Functions Triggers Rule Properties

+ ADD FIELD **DELETE SENSOR-TRAFFIC-COUNT**

source-ipv4-address
destination-ipv4-address
sensor-traffic-count
total-traffic-count

Field Name * ?
sensor-traffic-count

Description ?
Sensor octet count for IPv4 traffic measurement

Field Type
integer

Add to Rule Key ?

Ingest type (Field source)
Sensor

Sensor **Path *** **Data if missing**
 ipv4-flow-sensor octetDeltaCount Zero suppression Default value

Where (filter using expression)
+ ADD EXPRESSION

9. Click the **+ Add Field** button again and enter the following parameters to configure the fourth field, *total-traffic-count*:

Sensors **Fields** Vectors Variables Functions Triggers Rule Properties

+ ADD FIELD **DELETE TOTAL-TRAFFIC-COUNT**

source-ipv4-address
destination-ipv4-address
sensor-traffic-count
total-traffic-count

Field Name * ?
total-traffic-count

Description ?
Periodic sum of packet count for IPv4 measured in device

Field Type
integer

Add to Rule Key ?

Ingest type (Field source)
Formula

Formula **Field ***
Sum \$sensor-traffic-count

Time range *
10s

10. Click the **+ Add Field** button again and enter the following parameters to configure the fifth field, *traffic-count-maximum*:

Sensors **Fields** Vectors Variables Functions Triggers Rule Properties

+ ADD FIELD

source-ipv4-address
destination-ipv4-address
sensor-traffic-count
total-traffic-count
traffic-count-maximum

DELETE TRAFFIC-COUNT-MAXIMUM

Field Name * [?]
traffic-count-maximum

Description [?]
Maximum total traffic count

Field Type
integer

Add to Rule Key [?]

Ingest type (Field source)
Constant

Constant value *
{{traffic-count-max}}

11. Click the **+ Add Field** button once more and enter the following parameters to configure the sixth field, *traffic-count-minimum*:

Sensors **Fields** Vectors Variables Functions Triggers Rule Properties

+ ADD FIELD

source-ipv4-address
destination-ipv4-address
sensor-traffic-count
total-traffic-count
traffic-count-maximum
traffic-count-minimum

DELETE TRAFFIC-COUNT-MINIMUM

Field Name * [?]
traffic-count-minimum

Description [?]
Minimum total traffic count

Field Type
integer

Add to Rule Key [?]

Ingest type (Field source)
Constant

Constant value *
{{traffic-count-min}}

12. As the last step for the fields configuration, set the field aggregation time-range value to 10s:

Field aggregation time-range: 10s

Sensors **Fields** Vectors Variables Functions Triggers Rule Properties

13. Now move to the Variables tab, click the **+ ADD VARIABLE** button and create the *traffic-count-max* and *traffic-count-min* variables that are the constants for the *traffic-count-maximum* and *traffic-count-minimum* fields, respectively.

Sensors Fields Vectors **Variables** Functions Triggers Rule Properties

+ ADD VARIABLE

traffic-count-max
traffic-count-min

DELETE TRAFFIC-COUNT-MAX

Variable name * [?] traffic-count-max

Default Value [?] 10000

Type * [?] Integer

Description [?]
Maximum traffic count threshold in PPS

NOTE: Only the definition for the *traffic-count-max* is shown graphically. Choose an appropriate **Default Value** when configuring both *traffic-count-max* and *traffic-count-min* variables. The value shown above is for testing purposes only and may not be appropriate for your network.

14. Now move to the Triggers tab, click the **+ Add trigger** button and enter the following parameters to configure a trigger called *traffic-measurement-trigger*.

Sensors Fields Vectors Variables Functions **Triggers** Rule Properties

+ ADD TRIGGER

DELETE TRAFFIC-MEASUREMENT-TRIGGER

traffic-measurement-trigger

Trigger Name * [?] traffic-measurement-trigger

Frequency [?] 90s

Disable alert deduplication

Term traffic-anomaly-gr

WHEN

Left operand	Operator	Right operand	All in time range
\$total-traffic-count	>	Traffic-count-maximum	Enter a time range

+ ADD CONDITION

THEN

Color ■

Message
Total traffic count is above normal. Current total traffic count is \$total-traffic-count.

traffic-measurement-trigger

Trigger Name * [?]

traffic-measurement-trigger

Frequency [?]

90s

Disable alert deduplication

Term traffic-abnormal-gr ✕

Term traffic-abnormal-ls ✕

WHEN

Left operand	Operator	Right operand	All in time range
\$total-traffic-count	<	\$traffic-count-minimum	Enter a time range ⊖

[+ ADD CONDITION](#)

THEN

Color

Message

Total traffic count is below normal. Current total traffic count is \$total-traffic-count.

Evaluate next term

traffic-measurement-trigger

Trigger Name * [?]

traffic-measurement-trigger

Frequency [?]

90s

Disable alert deduplication

Term traffic-abnormal-gr ✕

Term traffic-abnormal-ls ✕

Term default-term ✕

WHEN

[+ ADD CONDITION](#)

THEN

Color

Message

Total traffic count is normal. Current total traffic count is \$total-traffic-count.

Evaluate next term

15. At the upper right of the window, click the **Save & Deploy** button.

Usage Notes:

- **Sensor Tab:**
 - The sensor name *ipv4-flow-sensor* is user-defined

- The sensor type is flow
- The sensor uses the predefined template *hb-ipfix-ipv4-template*
- **Variables Tab:**
 - The variables *traffic-count-max* and *traffic-count-min* are statically configured integers. In this case the values represent Bytes per second
 - These values are referenced in fields *traffic-count-maximum* and *traffic-count-minimum* and provide a reference point to compare against the *total-traffic-count* field
- **Fields Tab:**
 - Six fields are defined; some fields are used in the trigger settings while one field is referenced within another field
 - The field names are user-defined fields (UDF)
 - Fields *source-ipv4-address*, *destination-ipv4-address*, and *sensor-traffic-count* are extracting information from the flow sensor input
 - Path values for these fields identify specific values from the NetFlow messages, using naming according to [IPFIX Information Elements](#)
 - Fields *source-ipv4-address* and *destination-ipv4-address* have the Add to rule key setting enabled, indicating that this field should be shown as a searchable key for this rule on the device health pages
 - Field *total-traffic-count* - sums the IPv4 packet count from the *sensor-traffic-count* field every 10 seconds
 - The fields *traffic-count-maximum* and *traffic-count-minimum* are simply fixed values; the values are derived from the variables defined above
 - Field *aggregation time-range* - typically set to a value higher (longer) than individual field time range settings with the aim of reducing the frequency of information being sent to the database
- **Triggers Tab:**
 - The trigger name *traffic-measurement-trigger* is user-defined.
 - *frequency 90s* - Paragon Insights compares traffic counts every 90 seconds
 - In the term *traffic-abnormal-gr*:
 - When *\$total-traffic-count* (the periodic count of incoming IPv4 traffic) is greater than *\$traffic-count-maximum* (2500 Bps), show red and the message: "Total traffic count is above normal. Current total traffic count is : *\$total-traffic-count*".

- In the term `traffic-abnormal-ls`:
 - When `$total-traffic-count` (the periodic count of incoming IPv4 traffic) is less than `$traffic-count-minimum` (500 Bps), show yellow and the message: "Total traffic count is below normal. Current total traffic count is : `$total-traffic-count`".
- In the term `default-term`:
- Otherwise, show green and the message: "Total traffic count is normal. Current total traffic count is : `$total-traffic-count`".

RELATED DOCUMENTATION

| [Configure NetFlow Settings | 455](#)

About the Frequency Profiles

IN THIS SECTION

- [Tasks You Can Perform | 469](#)

Frequency profiles are a central location in which sensor and rule time frequencies can be managed. To understand frequency profiles, consider the following.

When defining rules in Paragon Insights you can:

- Define multiple rules that use the same sensor
- Define different sensor frequencies for each of the rules
- Apply all of these rules to the same device group/devices

This creates complexity in rule application and frequency adjustments within the individual rules:

- A key, consisting of sensor-path for OpenConfig and Native GPB sensors, or the tuple of file and table for iAgent sensors is used to identify the specific rules.
- Paragon Insights takes the minimum defined frequency for that sensor from the applied rules and uses it to subscribe to, or fetch, data from the devices.

- This make it hard to identify what the data rate should be for that sensor. To do that, you would have to go through the all the applied rules.
- A change in the sensor frequency of an applied rule might not take effect as intended.

To address these complexities, Paragon Insights needed a common place from which to control these frequencies. Frequency profiles can be created that allow you to manage sensor and rule frequencies from a single location and then apply the profiles in various locations in Paragon Insights. Application of these profiles allows for persistent and repeatable behavior in regard to frequencies for rules, sensors, triggers, formulas, references, learning periods, and hold times.

A sensor profile consists of a profile name and two optional sections: the sensors section and the non-sensors section. In each section, an entry consists of a sensor or rule name and a frequency. Frequency profiles are applied to device groups or network groups.

Tasks You Can Perform

To access the page from where you can create frequency profiles, go to **Configuration > Sensor > Settings**. You can perform the following tasks:

- Manage frequency profile. For more information, see "[Manage Frequency Profiles](#)" on page 469

RELATED DOCUMENTATION

| [Configure Offset Time](#) | 474

Manage Frequency Profiles

IN THIS SECTION

- [Configure a Frequency Profile](#) | 470
- [Edit a Frequency Profile](#) | 470
- [Clone a Frequency Profile](#) | 471
- [Delete a Frequency Profile](#) | 471

Frequency profiles are configured and managed in the Paragon Automation GUI. In the GUI, they are managed by navigating to the **Configuration > Data Ingest > Settings** page and selecting the **Frequency Profile** tab from the left side of the page.

NOTE: While the sections of the frequency profile are both optional, at least one section must be filled out per frequency profile if you want the applied profile to be able to do anything.

Configure a Frequency Profile

To configure a frequency profile:

1. Click **Configuration > Data Ingest > Settings** in the left-nav bar.
2. Click the **add (+)** icon to add a profile.
The Add Frequency Profile window appears.
3. Give the profile a name such as Profile1
If you configured rules with sensors, follow steps 3 through 5. If you configured rules without sensors, follow steps 6 through 8.
4. Click the **add (+)** icon to add sensors.
5. In the **Sensor Name** field, enter the sensor name as per the following guidelines:
 - *OpenConfig Sensors*: Enter the OpenConfig path for the desired sensor, such as /components or /interfaces.
 - *iAgent Sensors*: Enter the table name used in the sensor definition, such as ChassisAlarmTable or REutilizationTable
 - *SNMP*: Enter the sensor name such as npr_qmon_ext
 - *BYOI*: Enter <topic-name/ rule-name/ sensor-name>, such as topic1/rule1/sensor1
6. In the **Frequency** field, enter the appropriate frequency, such as 30seconds, 1minute, 2hours, and so on.
7. (Optional) Click the **add (+)** icon to add non-sensors.
8. In the **Rule Name** field, enter the rule name such as check-chassis-alarms.
9. In the **Frequency** field, enter the appropriate frequency, such as 45seconds, 3minutes, 1hour, and so on.
Repeat steps 3 through 5 or 6 through 8 as desired for the profile.
10. Click the **SAVE & DEPLOY** button to save and deploy the profile.
The new sensor profile is added to the list.

Edit a Frequency Profile

To edit a frequency profile:

1. Click **Configuration > Data Ingest > Settings** in the left-nav bar.

The Ingest Settings page is displayed.

2. Click the **Frequency Profile** tab to view the Frequency Profile page.
3. Select the *<Profile Name>* that you want to modify.
4. Click the **edit (pencil)** icon to edit the profile.

The Edit Frequency Profile page appears, displaying the same fields that are presented when you add a frequency profile.

5. Modify the parameters as shown in "[Configure a Frequency Profile](#)" on page 470.
6. Click the **SAVE** button to save the profile for later deployment or the **SAVE & DEPLOY** button to save and deploy immediately.

Clone a Frequency Profile

To clone an existing frequency profile:

1. Click **Configuration > Data Ingest > Settings** in the left-nav bar.

The Ingest Settings page is displayed.

2. Click the **Frequency Profile** tab to view the Frequency Profile page.
3. Select the *<Profile Name>* that you want to clone and then click the **Clone** button at the top-right corner of the page.

The Clone Frequency Profile page appears.

4. Specify an appropriate name for the cloned profile.
5. Click **OK** to save your changes.

A clone of the profile is created and listed on the Frequency profile page.

Delete a Frequency Profile

To delete an existing frequency profile:

1. Click **Configuration > Data Ingest > Settings** from the left-nav bar.


The Ingest Settings page is displayed.

2. Click the **Frequency Profile** tab to view the Frequency Profile page.
3. Select the frequency profile that you want to delete, and click the **delete (trash can)** icon.

The **CONFIRM DELETE FREQUENCY PROFILE** pop-up appears.

4. Do any one of the following:

Figure 31: Confirm Delete Frequency Profile Pop-up



- Click **Yes** to delete the frequency profile from the database. However, the changes are not applied to the ingest service.

NOTE:

- We recommended that you do not delete a frequency profile that is currently in use.
- After you delete a particular frequency profile from the database, you cannot apply that frequency profile to another device groups even if you have not deployed changes.
- You can also deploy changes to the ingest service or roll back the changes that you have already deleted, from the **Health Configuration Deployment Status** page. For more information, see "[Commit or Roll Back Configuration Changes in Paragon Insights](#)" on [page 185](#).

- Select the **Deploy changes** check box and then click **Yes** to delete the frequency profile from the database, and to apply the changes to the ingest service.
- (Optional) Click **No** to cancel this operation.

The Frequency Profile is deleted.

Usage Notes:

- *Profile Entries*–Multiple entries can be configured in each sensor profile section.
- *Override*–Sensor or rule frequency defined within an applied frequency profile overrides those defined within the individual rule or sensor.

- *Order of Precedence*—If a sensor or rule is defined in multiple frequency profiles, each with different frequency settings, the minimum frequency value for the sensor or rule is used.

SEE ALSO

[Apply a Frequency Profile | 473](#)

[About the Frequency Profiles | 468](#)

Apply a Frequency Profile

Frequency profiles are applied to Paragon Insights device groups or network groups. When you create or edit a device or network group, you apply frequency profiles by selecting them in the **Advanced** section of the **Edit Device Group** page. The following steps describe how you can apply frequency profiles on an existing device group. You can also follow the same steps while creating a device group.

To apply a frequency profile on a device group:

1. Go to **Configuration > Device Groups**.
In the Device Group Configuration page, select a device group.
2. Click on the pencil icon to edit the selected device group.
3. Click on **Advanced** and select a frequency profile from the drop-down menu in the *Ingest Frequency Profiles* field.
4. Click **Save** to only save the changes and **Save and Deploy** to deploy the frequency profile for the rules applied on the device group.

BEST PRACTICE: It is strongly recommended that you only apply frequency profiles to rules that make use of the Offset Time Unit feature.

The following steps describe how you can apply frequency profiles on an existing network group. You can also follow the same steps while creating a network group.

To apply a frequency profile on a network group:

1. Go to **Configuration > Network Group**.
In the Network Configuration page, select a network group.
2. Click on the pencil icon to edit the selected device group.
The Edit Network page appears.

3. Select a frequency profile from the drop-down menu in the *Ingest Frequency Profiles* field.
4. Click **Save** to only save the changes and **Save and Deploy** to deploy the frequency profile for the rules applied on the device group.

RELATED DOCUMENTATION

[Configure Offset Time | 474](#)

[Configure a Custom Rule in Paragon Automation GUI | 343](#)

Configure Offset Time

IN THIS SECTION

- [Offset Used in Formulae | 475](#)
- [Offset Used in Reference | 476](#)
- [Offset Used in Vectors | 477](#)
- [Offset Used in Trigger Term | 478](#)
- [Offset Used in Frequency Profile Applied to a Rule | 479](#)

The Paragon Insights offset time unit is used in conjunction with the "[About the Frequency Profiles](#)" on [page 468](#) to automatically manage time range calculations in various places within Paragon Insights rules. To understand the Paragon Insights offset function, consider the following scenario:

In Paragon Insights, you can define a rule which

- Uses a sensor to gather data with a frequency of 10 seconds
- Has a field that calculates the mean every 60 seconds

If you later decide to increase the frequency of the sensor to 60 seconds, then calculating the mean every 60 seconds would not make any sense. The result is that you would have to manually update the field calculation any time up want to change the sensor frequency.

You can set the time range for the mean calculation to a value of 60, or 60offset, rather than 60s, or 60 seconds. Using an offset time value rather than a static time value tells Paragon Insights to automatically

multiply the sensor frequency by the numeric value of the offset. Thus, any change to sensor frequency will automatically be included in the time range calculation for a formula.

An offset time unit can be used in place of standard time units in the following time range locations:

- Formulae
- References
- Vectors
- Trigger Frequency
- Trigger Term

Offset Used in Formulae

In this example, we are creating a rule, *Rule1*, with a sensor, *Sensor1*. The sensor frequency is set to 10 seconds.

In [Figure 32 on page 475](#) below, you can see the **Fields** block definition for *Rule1* with the **Sensors** block definition framed in green. The formula, *formula1*, has a **Time range** value of *2o*

Figure 32: Offset in a Formula

The screenshot displays the configuration page for a rule named 'Rule1'. The interface includes a top navigation bar with 'Rule: external / Rule1', 'Rule Frequency:', and buttons for 'SAVE & DEPLOY', 'SAVE', 'DELETE', and 'CLONE'. Below this is a 'Description' field containing 'Demonstration Rule' and a 'Synopsis' field with the text 'This rule is used only to demonstrate various rule features and concepts.' A 'Field aggregation time-range:' field is also present.

The main configuration area is divided into tabs: 'Sensors', 'Fields', 'Vectors', 'Variables', 'Functions', 'Triggers', and 'Rule Properties'. The 'Fields' tab is active, showing a list of fields on the left with 'formula1' and 'field1'. The main area for the 'Fields' block includes:

- Field Name ***: formula1
- Description**: Add a description for this field
- Field Type**: integer
- Add to Rule Key**: (toggle off)
- Ingest type (Field source)**: Formula
- Formula**: Max
- Field ***: \$field1
- Time range ***: 2o

The 'Sensors' tab is also visible, showing a list of sensors on the left with 'Sensor1'. The main area for the 'Sensors' block includes:

- ADD SENSOR** button
- Sensor1** (selected)
- Sensor Name ***: Sensor1
- Sensor Type**: Open Config
- Sensor Path ***: /interfaces
- Frequency ***: 10s

 A green rectangular box highlights the 'Sensors' block configuration area.

Usage Notes for Offset Used in Formulas

- An offset value applied to the time range of a formula multiplies the rule, sensor, or trigger frequency of that rule.
- The result of this example is that the time range for the Max formula is now 2 times the *Sensor1* frequency of 10 seconds, or 20 seconds ($2 * 10s = 20s$).
- If a frequency profile of 30 seconds is applied to the rule used in the example, then the resulting time-range would be 60 seconds ($2 * 30s = 60s$).
- Offset time can be applied to the following formulas: latest, count, min, max, sum, mean, on-change, and stdev.

Offset Used in Reference

In this example, there are two rules in play, but only one is shown. The unseen rule, Rule1, is in the topic routing-engines and is named routing-engines (routing-engines/routing-engines). Rule1 has a frequency of 20 seconds. Rule1 is referenced .

Rule2, shown in [Figure 33 on page 476](#), is a network rule which has a reference field named *ref1*. The field, *ref1*, references back to Rule1 through the **Reference XPath Expression** with a **Time Range** setting of 3offset.

Figure 33: Offset Time Used in Reference Field

The screenshot shows the configuration page for a Network Rule named 'Rule2'. The rule is associated with the 'external' sensor and has a frequency of '10s'. The description is 'Demonstration Rule' and the synopsis is 'This rule is used only to demonstrate various rule features and concepts.' The 'Field aggregation time-range' is currently empty.

The 'Fields' tab is active, showing a list of fields on the left: 'formula1' and 'ref1'. The 'ref1' field is selected and its configuration is shown on the right. The configuration includes:

- Field Name:** ref1
- Description:** Add a description for this field
- Field Type:** integer
- Add to Rule Key:** (unchecked)
- Ingest type (Field source):** Reference
- Reference XPath expression:** `"/device-group[device-group-name='Core4']/device[device-id='R1']/topic[topic-name='routing`
- Time Range:** 3offset
- Data if missing:** Default value

Buttons for '+ ADD FIELD', 'DELETE REF1', 'SAVE & DEPLOY', 'SAVE', 'DELETE', and 'CLONE' are visible at the top of the configuration area.

Usage Notes for Offset Used in References

- Offset values used in references multiply the frequency of the referenced rule by the offset value. In this case, the 20 second frequency of Rule1 is multiplied by 3, resulting in a 60 second time-range ($3 * 20s = 60s$).
- If a frequency profile of 60 seconds (60s) was applied to Rule1, the time range for the reference would increase to 180 seconds ($3 * 60s = 180s$).

Offset Used in Vectors

In this example, there are 3 rules at play. Rules *Rule1* and *Rule2* are not shown but are referenced by the vector.

Rule1 is in topic line-cards and is named line-cards (line-cards/line-cards) and has a frequency of 20 seconds (20s).

Rule2 is in topic routing-engines and is named routing-engines (routing-engines/routing-engines) and has a frequency of 30 seconds (30s).

In *Rule3* below, we have defined a vector, *vector1* that consists of 2 path references and 1 field. The vector has a **Time Range** defined as 3offset. [Figure 34 on page 477](#) shows the vector block definition in the Paragon Insights GUI.

Figure 34: Offset Time Used in Vector Block

The screenshot shows the configuration interface for a Network Rule. At the top, the rule is named 'Rule3' with a frequency of '10s'. Below this, there are tabs for 'Sensors', 'Fields', 'Vectors', 'Variables', 'Functions', 'Triggers', and 'Rule Properties'. The 'Vectors' tab is active, showing a list of vectors with 'vector1' selected. The configuration for 'vector1' includes:

- Vector Name:** vector1
- Ingest Type:** path
- Time-Range:** 3offset
- References:**
 - "/device-group[device-group-name='Core4']/device[device-id='R1']/topic[topic-name='routing-engines']/rule[rule-name='routing-engines']/field[slot='0']/ref1" x
 - "/device-group[device-group-name='Core4']/device[device-id='R1']/topic[topic-name='line-cards']/rule[rule-name='line-cards']/memory" x
 - \$formula1 x

Usage Notes for Offset Used in Vectors

- An offset value defined in the time range of a vector multiplies the sensor or rule frequency of the referenced rule or sensor. If a field from the same rule is used as the reference, then the frequency of the rule containing the vector is used.
- When multiple references or fields are defined for a vector and offset time is used, the offset value is applied to each path independently. So, for this example in which our offset value is 3 (3offset):
 - The path reference to Rule1: `"/device-group[device-group-name='Core4']/device[device-id='R1']/topic[topic-name='line-cards']/rule[rule-name='line-cards']/memory"` which has a frequency of 20 seconds, would result in a time range of 60 seconds for the vector ($3 * 20s = 60s$).
 - The path reference to Rule2: `"/device-group[device-group-name='Core4']/device[device-id='R1']/topic[topic-name='routing-engines']/rule[rule-name='routing-engines']/field[slot='0']/ref1"` which has a frequency of 30 seconds, would result in a time range of 90 seconds for the vector ($3 * 30s = 90s$).

NOTE: There are no spaces or line breaks in the path references. They are added to this document only to enhance readability.

- The field reference to `formula1`: Since `formula1` is a normal field used within the vector block, the rule frequency of the current rule is used as a basis. So, the time range for `formula1` is 30 seconds ($3 * 10s = 30s$).
- If a frequency profile of 60 seconds is applied to Rule1 (line-cards/line-cards), then the time range for that path in the vector would be 180 seconds ($3 * 60s = 180s$).
- If a frequency profile of 120 seconds is applied to Rule2 (routing-engines/routing-engines), then the time range for that path in the vector would be 360 seconds ($3 * 120s = 360s$).
- The time range for the field reference, `formula1`, would remain the same as before at 30 seconds unless a change is made to Rule3 or a frequency profile with a different time is applied to Rule3.

Offset Used in Trigger Term

In this example, we have one rule, Rule1. [Figure 35 on page 479](#) below shows the trigger block definition with the sensor block definition overlaid with a green border.

The rule, *Rule1*, has a sensor frequency of 10 seconds (10s) applied.

The trigger itself, *Trigger1*, has an offset frequency of 2 (2o) applied.

The trigger term, *term1*, has its own time range offset of 2 (2o) applied as well.

Figure 35: Offset Time Used in Trigger Block

The screenshot displays the configuration for a rule named 'Rule1' in the 'external' namespace. The rule is a 'Demonstration Rule' with a synopsis: 'This rule is used only to demonstrate various rule features and concepts.' The configuration is divided into several tabs: Sensors, Fields, Vectors, Variables, Functions, Triggers, and Rule Properties.

Triggers Configuration:

- Trigger Name:** trigger1
- Frequency:** 2o
- Disable alarm deduplication:**
- Term:** term1
- WHEN:**
 - Left operand: \$name
 - Operator: FPC1
 - Operator: All in time range
 - Frequency: 2o

Sensors Configuration:

- Sensor Name:** Sensor1
- Sensor Type:** Open Config
- Sensor Path:** /interfaces
- Frequency:** 10s

Usage Notes for Offset Time Used in Triggers

- An offset value defined in trigger frequency multiplies the sensor or rule frequency. So, the offset value of 2 in *trigger1* causes the trigger frequency to be interpreted as 20 seconds ($2 * 10s = 20s$) because the sensor frequency of the rule is used as the basis.
- An offset value defined in a trigger term multiplies the trigger frequency value. So, the offset value of 2 in *term1* causes the term frequency to be interpreted as 40 seconds ($2 * 20s = 40s$).

Offset Used in Frequency Profile Applied to a Rule

In this example, we have 2 rules in the topic external, and one frequency profile, *prof1*.

The rules are:

- *external/test*: The *test* rule has a sensor named components which is an OpenConfig sensor with a sensor path of */components* and a sensor frequency of 10 seconds.

It also has a trigger, *trig1* with a frequency of 2o or 2offset. The trigger has a term named *Term_1*, which is not used in the example.

The screenshot displays a configuration interface for a rule named `external / test`. The interface is divided into several sections:

- Rule Information:** Includes fields for `Rule:` (external / test), `Description:` (Description), `Synopsis:` (Synopsis), and `Field aggregation time-range:`.
- Sensors Section:** Features a tabbed interface with `Sensors`, `Fields`, `Vectors`, and `Variables`. It includes an `+ ADD SENSOR` button and a sensor named `components`. The configuration for `components` shows:
 - Sensor Name ***: components
 - Sensor Type**: Open Config
 - Sensor Path ***: /components
 - Frequency ***: 10s
- Triggers Section:** Features a tabbed interface with `Sensors`, `Fields`, `Vectors`, `Variables`, `Functions`, and `Triggers`. It includes an `+ ADD TRIGGER` button and a trigger named `trig1`. The configuration for `trig1` shows:
 - Trigger Name ***: trig1
 - Frequency ?**: 2o
 - Disable alarm deduplication**:
 - Term**: Term_1

- *external/ref*: The *ref* rule is a non-sensor rule, which means that the rule uses a reference field, *trigger_reference*, to reference the sensor defined in another rule; in this case *external/test*.

Rule: external / ref Rule Frequency: 30s Network Rule SAVE & DEPLOY SAVE

Description: Description

Synopsis: Synopsis

Field aggregation time-range:

Sensors **Fields** Vectors Variables Functions Triggers Rule Properties

+ ADD FIELD

trigger_reference

Field Name * (?)

trigger_reference

Description (?)

Add a description for this field

Field Type

Field type

Add to Rule Key (?)

Ingest type (Field source)

Reference

Reference XPATH expression *

/topic[topic-name='external']/rule[rule-name='test']/trigger[trigger-name=trig1]/color

Time Range

2offset

And the frequency profile is:

- *prof1*: The *prof1* profile sets the frequency for the */components* sensor at 30 seconds.

Edit Frequency Profile: prof1

Name

SENSORS

Sensor Name	Frequency *
<input type="text" value="/components"/>	<input type="text" value="30seconds"/> -

+ ADD SENSORS

NON-SENSORS

+ ADD NON-SENSORS

CANCEL
SAVE
SAVE & DEPLOY

Usage Notes for Offset Time Used in Trigger Reference

- An offset value defined in trigger frequency multiplies the sensor or rule frequency. So, the offset value of 2 in *trig1* causes the trigger frequency to be interpreted as 20 seconds ($2 * 10s = 20s$) because the sensor frequency of the rule is used as the basis.
- An offset value defined in a rule reference *external/ref* is 2. This value multiplies the frequency value of 30s in frequency profile *prof1* (the frequency configure in frequency profile *prof1* (30s) takes precedence over the sensor frequency of 10s configured in *external/test/rule*). So, the offset value of 2 in rule *ref* causes the frequency to be interpreted as 60 seconds ($2 * 30s = 60s$).

RELATED DOCUMENTATION

[Manage Frequency Profiles](#) | 469

Configure a Rule Using Server Monitoring Sensor

In the following example, you can use server monitoring sensor to collect disk read data from servers. You can configure fields for total disk read size, time taken to perform the reads, and name of the device that has the disk (key field). You can also calculate rate of disk read and configure a trigger alert when the total disk read exceeds a preset threshold.

1. Click **Configuration > Rules** .

2. On the Rules page, click **+ Add Rule**.
3. Enter topic name as **server.monitoring** and set the rule name after the slash ('/'). In this example, rule name can be check-disk-read.
The rule name in the top row of the rule page follows the topic/rule name format. The default topic name is 'external' when you add a new rule.
4. Add a description and synopsis for your rule.
5. Click **+ Add Sensor** and enter a name for the sensor. For example, disk.
6. Select **Server Monitoring** as sensor type.
7. Enter sensor path as **/node/disk**.
If you add a / at the end of the path, you get sensor paths for disk reads, writes, and written records.
8. Enter a Frequency (in seconds) for the sensor. For example, 30s.
The minimum usable sensor frequency is 15 seconds. It takes at least 15 seconds before you see data from the ingest.
9. Go to Fields tab and click **+ Add Field** and enter the Field Name as **device-name**.
This field collects the name of the device containing the disk that generates read data.
10. Select Field Type as **string**.
11. Enable **Add to Rule Key**.
12. Select Ingest Type (Field Source) as **Sensor**.
13. Select the name of the sensor in the Sensor field. In this example, select disk.
14. Type **Device** in the Path field.
15. Go to Fields tab and click **+ Add Field** and enter the Field Name as **disk-read-total**.
This field collects the total size of disk reads in bytes.
16. Select Field Type as **float** and Ingest Type (Field Source) as **Sensor**.
17. Select the name of the sensor for the Sensor field. In this example, select **disk**.
18. Select Path as **/node/disk/read/bytes/total**
19. Click **+ Add Field** and enter the Field Name as **disk-read-rate**.
This field calculates the rate of change using the field value in disk-read-total.
20. Select Field Type as **float** and Ingest Type (Field Source) as **Formula**.

21. In Formula, select **Rate of Change**.
22. In Field, select the field name **disk-read-total**.
23. Click **+ Add Field** and enter the Field Name as **read-threshold**.
This field contains a constant value for disk read threshold.
24. Select Field Type as **float** and Ingest Type (Field Source) as **Constant**.
25. In Constant Value, enter a threshold value for disk reads. For example, **5**.
26. Go to Triggers tab and click **+Add Trigger**.
27. Enter a name for the trigger such as **read-trigger**.
28. Enter a Frequency value. For example, **2o** (2 offset).
If you set frequency as **2o** or **2 offset**, it multiplies the static frequency you set for sensor frequency by 2.
29. Click **+Add Term** and enter a term name. For example, **high-disk-usage**.
30. In the When statement, select left operand as **disk-read-total** field, right operand as **read-threshold** field, and the operator as **Increase At Least by Value**.
31. In the Then statement, set red as the color and enter Message as **high disk read value**.
32. Do one of the following:
 - Click **Save** to save the rule configuration. Paragon Automation does not deploy the rule configuration.
 - Click **Save and Deploy** to deploy the configuration in Paragon Automation Platform.

To collect data on metrics, you must add the rule to a Playbook. Then, apply a Playbook instance to the device or the network groups.

When you start collecting server metrics, you can see the logs by providing the pod IDs for the ingest.

- Log in to the Paragon Automation management CLI.
- Type the command `/var/local/healthbot/healthbot k logs server-monitoring pod id`.

RELATED DOCUMENTATION

[Server Monitoring Sensor | 440](#)

[Create and Run a Playbook Instance | 314](#)

Configure Native GPB Ingest

The Junos device must run Junos OS Version: 15.1 or later to use the Native GPB sensor type on a Junos OS device.

To configure a Junos OS device to use Native GPB, you must configure a sensor profile, a streaming profile, an export profile, and a resource string.

- *Streaming profile* – The profile associated with the server that collects exported data streamed by a monitored system resource. You can configure more than one streaming server. To collect data, you must associate a configured server with one or more configured sensors.

To configure the server that collects data, you must also configure a destination IP address and a destination port on the Junos device. The same port must be configured at the remote end, in the Paragon Automation Platform.

- *Export profile* – An export profile defines the parameters of the export process of data generated through the sensor. You can associate multiple sensor profiles with an export profile but each sensor profile can be associated with only one export profile.
- *Sensor profile* – A sensor profile defines the parameters of the system resource through resource strings (such as `/junos/services/ldp/label-switched-path/transit/usage/`) that is monitored.

See [Configuring a Junos Telemetry Interface Sensor](#) for more information on configuring the individual profile.

The following is sample configurations of streaming server profile, export profile, and sensor profiles.

```
##Streaming Server Profile
set services analytics streaming-server COLLECTOR-1 remote-address <virtual-IP-address-Insights-
Services>
set services analytics streaming-server COLLECTOR-1 remote-port 22000
##Export Profile
set services analytics export-profile EXP-PROF-1 local-address <local-router-IP>
set services analytics export-profile EXP-PROF-1 local-port 22001
set services analytics export-profile EXP-PROF-1 reporting-rate 30
set services analytics export-profile EXP-PROF-1 format gpb
set services analytics export-profile EXP-PROF-1 transport udp
##Sensor Profile
set services analytics sensor SENSOR-1 server-name COLLECTOR-1
set services analytics sensor SENSOR-1 export-name EXP-PROF-1
set services analytics sensor SENSOR-1 resource <resource> # example /junos/system/linecard/
interface/
```

NOTE: The virtual IP address for Insights Services (configured during Paragon Automation Platform installation) is used for Native GPB, Syslog, and SNMP ingest configurations.

To configure streaming server port in Paragon Insights GUI:

1. Go to **Configuration > Data Ingest > Settings**.
2. Select **Native GPB** tab on the Ingest Settings page.
3. Enter the port number. The port number must be the same as the remote-port configured in the streaming server profile.
You can use the toggle button to enable or disable the **Port** field.
4. Click **Save & Deploy** to enable the sensor to collect data in your network.

RELATED DOCUMENTATION

| [Sensors Overview](#) | 429

Configure sFlow Settings

IN THIS SECTION

- [Configure Devices to Send sFlow Packets](#) | 487
- [Configure sFlow Ingest](#) | 488
- [Delete sFlow Settings](#) | 493
- [Configure sFlow in Devices and Device Groups](#) | 496
- [Configure a Rule Using sFlow](#) | 498

This section describes the configuration of sFlow ingest and configurations in device or device group configuration to stream sFlow packets in Paragon Automation.

Configure Devices to Send sFlow Packets

When you configure a device to send sFlow to a collector, you simply set a source IP address (IP address of the collector), sample-rate, polling interval, UDP port, and interface to capture from. There is no opportunity to filter or choose what data gets sent from the device side.

NOTE: The IP address of collector is the virtual IP address of Paragon Insights services you set while installing Paragon Automation Platform.

The following is an example configuration snippet to configure an MX series router to send sFlow packets.

```
[edit protocols]
  set sflow collector 10.234.32.46 udp-port 5600

  set sflow interfaces ge-0/0/0
  set sflow polling interval 20
  set sflow sample-rate egress 1000
  set sflow interfaces ge-0/0/1 polling-interval 10 sample-rate ingress 1000
```

The following is an example configuration snippet to configure an EX series switch to send sFlow packets.

```
[edit protocols]
  set sflow collector 10.234.32.46 udp-port 5600

  set sflow interfaces ge-0/0/0
  set sflow polling interval 20
  set sflow sample-rate egress 1000
```

The following example shows the output from a switch already configured to send sFlow packets to a collector at IP address 10.204.32.46.

```
[edit protocols sflow]
  user@switch# show
  polling-interval 20;
  sample-rate egress 1000;
  collector 10.204.32.46
  {
```

```

    udp-port 5600;
  }
  interfaces ge-0/0/0.0;

```

Configure sFlow Ingest

As with other ingest methods, navigate to **Configuration > Data Ingest > Settings** and choose the **sFlow** tab.

The **sflow Settings** are broken down into 4 sections:

- **Sample** There are two pre-defined sample categories and each is represented in the sFlow header as an integer sample-type value. [Table 78 on page 488](#) below shows the sample types and their numeric value.

Table 78: sFlow Sample Types

Sample Type	Integer Value in sFlow Header
counter-sample	2
expanded-counter-sample	4
flow-sample	1
expanded-flow-sample	3

NOTE: The difference between the expanded sensor types and the non-expanded sample types is the size of the data fields. The field names and types are the same, but the field sizes are larger in the expanded sample types.

Packet definitions for these sample types can be found here: [sFlow Samples](#)

[Table 79 on page 489](#) shows the other fields contained in an sFlow sample header (by sample type) along with the field type.

Table 79: sFlow Packet Header Fields

field type/size in bits	counter-sample	flow-sample
integer/32	sampleSequenceNumber	sampleSequenceNumber
integer/8	sourceIDType <ul style="list-style-type: none"> • 0 = SNMP interface index • 1 = VLAN ID (smonVlanDataSource) • 2 = Physical entity (entPhysicalEntry) 	sourceIDType <ul style="list-style-type: none"> • 0 = SNMP interface index • 1 = VLAN ID (smonVlanDataSource) • 2 = Physical entity (entPhysicalEntry)
integer/24	sourceIDValue	sourceIDValue
integer/32	n (the number of sampled records contained in the Counter sample)	sampleSamplingRate
integer/32	-	samplePool (number of packets that could have been sampled)
integer/32	-	sampleDroppedPackets (number of packets dropped due to lack of resources)
integer/8	-	sampleInputInterfaceFormat (input interface type)
integer/32	-	sampleInputInterfaceValue (input interface (SNMP interface index))
integer/1	-	sampleOutputInterfaceFormat (output interface type)

Table 79: sFlow Packet Header Fields (Continued)

field type/size in bits	counter-sample	flow-sample
integer/33	-	sampleOutputInterfaceValue (SNMP interface index)
integer/32	-	n (the number of flow records)
data	counter records	flow records

- Flow Record** The **Flow Record** section provides the tools needed to define the different types of flow that might be seen in an sFlow capture. Paragon Automation ships with 16 types of pre-defined flow records, each of which have a format number and a sensor path for use in defining sFlow rules, shown in [Table 80 on page 490](#) below. There are several fields in each type of flow record. These can be seen by selecting the desired record type from the list and clicking the **edit (pencil)** button.

Table 80: Flow Record Types

Record Type	Format Number	Sensor Path Value
raw packet headers	1	/sflow-v5/flow-sample/raw-packet-header
Ethernet frame data	2	/sflow-v5/flow-sample/ethernet-frame-data
IPv4 data	3	/sflow-v5/flow-sample/ipv4-data
IPv6 data	4	/sflow-v5/flow-sample/ipv6-data
extended switch data	1001	/sflow-v5/flow-sample/extended-switch-data

Table 80: Flow Record Types (Continued)

Record Type	Format Number	Sensor Path Value
extended router data	1002	/sflow-v5/flow-sample/extended-router-data
extended gateway data	1003	/sflow-v5/flow-sample/extended-gateway-data
extended user data	1004	/sflow-v5/flow-sample/extended-user-data
extended URL data	1005	/sflow-v5/flow-sample/extended-url-data
extended MPLS data	1006	/sflow-v5/flow-sample/extended-mpls-data
extended NAT data	1007	sflow-v5/flow-sample/extended-nat-data
extended MPLS tunnel	1008	/sflow-v5/flow-sample/extended-mpls-tunnel
extended MPLS VC	1009	/sflow-v5/flow-sample/extended-mpls-vc
extended MPLS FEC	1010	/sflow-v5/flow-sample/extended-mpls-fec
extended LVP FEC	1011	/sflow-v5/flow-sample/extended-mpls-lvp-fec
extended VLAN tunnel	1012	/sflow-v5/flow-sample/extended-vlan-tunnel

When you configure rules for sFlow, you can choose from any of these record types. You can create new flow records by clicking the **add (+)** icon on the **Sflow Settings** page.

- **Counter Record** The **Counter Record** section provides the definition for the two pre-defined counter record types. There are two types of counter records, ethernet-interface-counters and generic-interface-counters. Generic interface counters are format number 1 and Ethernet interface counters are format number 2. The sensor path for generic interface counters is /sflow-v5/counter-sample/generic-interface-counter. The sensor path for Ethernet interface counters is /sflow-v5/counter-sample/ethernet-interface-counter.

The fields available within the counter records are the possible errors and the countable statistics such as:

- frame errors
- collisions
- deferred transmissions
- transmit errors
- administration status
- operational status
- input packets
- output packets
- input errors
- output errors
- and others

You can use either the generic interface counter or Ethernet interface counter in rules that you define. The counter sensors can be defined to pick even single fields from either of the available counters. You can create additional counter record types by clicking the **add (+)** icon on the **Sflow Settings** page (**Counter Record** section).

- **Protocol** The **Protocol** section provides a means to define which protocol the sFlow captures contain and allow for the decoding of many network protocols. The fields that are contained in each protocol entry are the same fields as would be seen in a frame or packet of that type. For example, an Ethernet frame would have a destination MAC address, a source MAC address, and an ethernet-next-header-type field. The fields defined in any

protocol you want to decode must appear in the protocol definition in the same order as they would appear in the packet or frame.

The number column that appears is the IANA protocol number assigned to that protocol. For example, the tcp protocol is protocol number 6.

NOTE: On the **Sample**, **Flow Record**, and **Counter Record** sections, there is an **Enterprise** column. This column is for the use of vendor-specific or custom decoding details. For example, a Foundry ACL-based flow sample has the enterprise value 1991, Format 1, includes additional fields specifically for that Foundry flow. In most instances, the Enterprise value is 0.


Delete sFlow Settings

To delete sFlow settings:

1. Click **Configuration > Data Ingest > Settings** from the left-nav bar.
The Ingest Settings page is displayed.
2. Click the **sFlow** tab to view the **sFlow Settings** page.
3. Do one of the following:
 - To delete an sFlow sample:
 - a. Click **Sample** to view the list of sFlow samples.
 - b. Select the sFlow sample that you want to delete.
 - c. Click the **delete (trash can)** icon.

The **CONFIRM DELETE SAMPLE** pop-up appears.


Figure 36: Confirm Delete Sample Pop-up



- To delete a flow record:
 - a. Click **Flow Record** to view the list of flow records.
 - b. Select the flow record that you want to delete.
 - c. Click the **delete (trash can)** icon.

The **CONFIRM DELETE FLOW RECORD** pop-up appears.

Figure 37: Confirm Delete Flow Record Pop-up




- To delete a counter record:

- a. Click **Counter Record** to view the list of counter records.
- b. Select the counter record that you want to delete.
- c. Click the **delete (trash can)** icon.

The **CONFIRM DELETE COUNTER RECORD** pop-up appears.


Figure 38: Confirm Delete Counter Pop-up



- To delete a protocol:
 - a. Click **Protocol** to view the list of protocols.
 - b. Select the protocol that you want to delete.
 - c. Click the **delete (trash can)** icon.

The **CONFIRM DELETE PROTOCOL** pop-up appears.

Figure 39: Confirm Delete Protocol Pop-up



4. In the pop-up that appears, do one of the following:

- Click **Yes** to delete the sFlow setting from the database. However, the changes are not applied to the ingest service.

NOTE:

- We recommended that you do not delete an sFlow setting that is currently in use.
- After you delete an sFlow setting from the database, you cannot configure that sFlow setting in new devices or device groups even if you have not deployed changes.
- You can also deploy changes to the ingest service or roll back the changes that you have already deleted, from the **Health Configuration Deployment Status** page. For more information, see "[Commit or Roll Back Configuration Changes in Paragon Insights](#)" on [page 185](#).

- Select the **Deploy changes** check box and then click **Yes** to delete the sFlow settings from the database, and to apply the changes to the ingest service.
- (Optional) Click **No** to cancel this operation.

The sFlow Setting is deleted.

Configure sFlow in Devices and Device Groups

To process sFlow packets, Paragon Automation depends on rule configuration. It also requires that you enable sFlow in the device group and device definition. This section describes sFlow enablement, and rule and sensor configuration options for sFlow.

First, to enable sFlow, you must enter at least one IP address in the device definition under **Flow Source IPs**, and enter at least one port number in the device group definition under **sFlow Ports**. [Figure 40 on page 497](#) below is a composite image that shows the device definition overlaid with the device group definition. The appropriate sections of each window are highlighted in red.

Figure 40: Enable sFlow Composite Image

The image shows two overlapping configuration windows. The background window is titled "Edit Device Group" and has a sidebar with sections: Authentication, Protocols, Device ID Details, and Advanced. The "Ports" section is expanded, showing fields for Native Ports, Flow Ports, sFlow Ports, Syslog Ports, and SNMP Notification Ports, each with a "Comma separated values" input field. The "sFlow Ports" field is highlighted with a red box. The foreground window is titled "Edit 'vmx102'" and has a sidebar with sections: Authentication, Protocols, Device ID Details, and Advanced. The "Device ID Details" section is expanded, showing fields for System ID to use for JTI, Management IP, Flow Source IPs, Syslog Source IPs, Syslog Host Names, and SNMP Source IPs. The "Flow Source IPs" field is highlighted with a red box. At the bottom of the foreground window, there are "Cancel" and "OK" buttons.

The devices in the group send their sFlow packets to Paragon Automation over the configured UDP port from the configured IP address(es). The port number(s) used in these definitions must be unique across the entire Paragon Automation installation.

NOTE:

- The one or more **Flow Source IPs** that you enter must match an IP address that can be mapped from the **Hostname/IP Address/Range** field in the device definition. If devices send sFlow packets, but Paragon Automation cannot match the source IP to a defined device IP, then the packets are dropped without decoding.
- Paragon Automation cannot differentiate sFlow from NetFlow by looking at the packets. If you are using both NetFlow and sFlow, the port numbers must also be unique between the two flow types.

Due to the nature of sFlow and the potentially huge amount of data that can come from even a single device, we recommend the following best-practices for managing sFlow ingest:

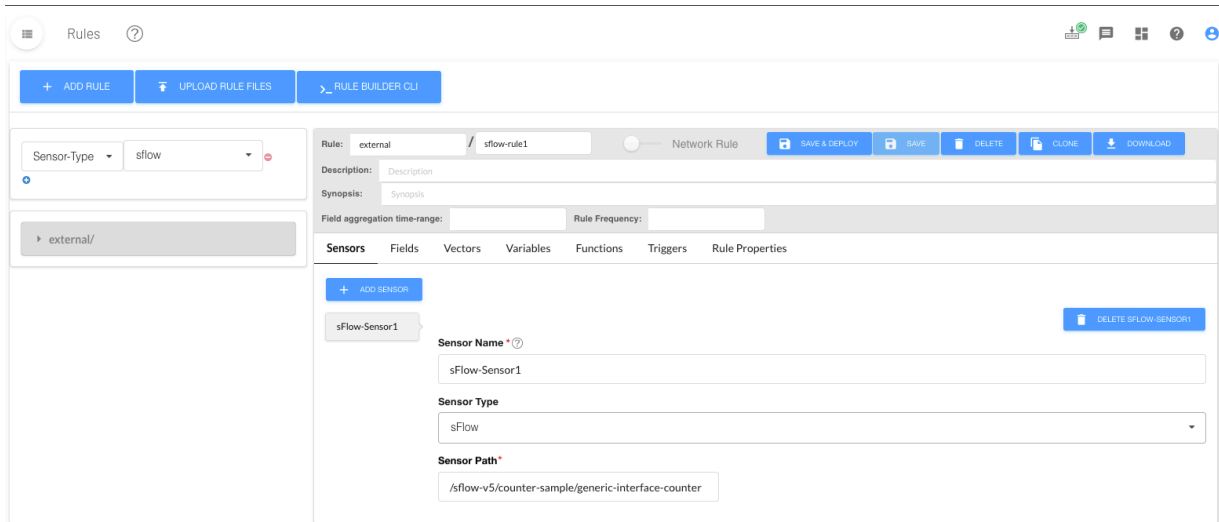
BEST PRACTICE:

- Use unique ports from the range: UDP/49152 to UDP/65535 for sFlow.
- Use periodic aggregation to reduce the number of write procedures in the TSDB.
- Do not enable the raw table data storage option in sFlow unless sufficient high-speed storage is available for Paragon Insights TSDB.

Configure a Rule Using sFlow

As with other rule definitions, sFlow rules are made up of sensors, fields, vectors, and so on. An sFlow sensor has a **Sensor Name**, a **Sensor Type** of sFlow, and an **sFlow Path** as shown in [Figure 41 on page 499](#).

Figure 41: sFlow Sensor Definition



The sensor path serves a big role in sensor definition. Paragon Automation uses the sensor path to define not only the sFlow flow type, but the sample type, record type, protocol, and other custom path elements if needed.

RELATED DOCUMENTATION

[Configure NetFlow Settings | 455](#)

Configure SNMP Ingest

Paragon Insights supports three methods of collecting telemetry data using SNMP. The ingest, also known as request-response, is a pull mode method in which Paragon Insights requests for telemetry data from the devices. The trap and inform notifications are push mode methods in which the devices notify Paragon Insights about key performance indicator events that prevents the devices from functioning as expected.

Paragon Insights supports SNMPv3 alongside the current SNMPv2c as an ingest method. Users with sp-admin role can select any version of SNMP in the Paragon Insights GUI.

SNMPv3 ingest provides you with an option to set authentication and privacy credentials to leverage the following features:

- Authentication – Identifies and verifies the origin of an SNMPv3 message.

- Privacy – Prevents packet analyzers from snooping the content of messages by encrypting them.
- Integrity – Ensures that the content of SNMP messages is not altered while in transit without authorization.

Table 81 on page 500 lists the supported authentication and privacy algorithms in SNMPv3 protocol.

Table 81: Authentication and Privacy Algorithms

Feature	Algorithm
Supported authentication algorithms	MD5
	SHA-1
Supported privacy algorithms	DES
	3DES
	AES

In Paragon Insights, SNMPv3 ingest can be set at the device level. The configuration of SNMPv3 and SNMPv2c is mutually exclusive.

NOTE: If a device is not configured for SNMP ingest, Paragon Insights uses SNMP v2c with **SNMP Community** set to *public* as the default settings.

NOTE: In Paragon Insights, the SNMPv2c and SNMPv3 ingest and trap configurations share the same workflow.

To configure SNMP ingest at the device level:

1. Click the **Configuration > Device** option in the left navigation bar.
2. Select a device by clicking on the checkbox and click the edit device button (pencil icon).
The **Edit *Device-Name*** window appears.
3. Click on **Protocol > SNMP**.
4. Enter the necessary values in the text boxes and select the appropriate options for the device.

The following table describes the attributes in the **Edit *Device-Name*** window:

Attributes	Description
SNMP	
Version	Select either v2c or v3 in the drop-down menu.
Get Community (Only for SNMPv2c)	<p>Enter an SNMP Community string for SNMPv2c ingest.</p> <p>In SNMPv2c, Community string is used to verify the authenticity of the ingest (request-response) message issued by the SNMP agent (devices such as routers, switches, servers, and so on).</p>
Port	Port number required for SNMP ingest (request-response) messages. The standard port number is 161 .
Timeout	<p>Enter the timeout period in seconds between 0 and 65535 for SNMP notifications.</p> <p>Timeout denotes the number of seconds after which the SNMP agent stops re-sending notifications.</p>
Retry Count	<p>Enter a retry count between 0 and 255.</p> <p>Retry count denotes the number of times the SNMP agent re-transmits the SNMP notification.</p>
V3 Username	Enter a username for SNMPv3 ingest (request-response), trap, and inform notifications.

(Continued)

Attributes	Description
V3 Context Name	<p>(Optional) Enter a context name for SNMPv3 trap and inform notifications.</p> <p>Context in SNMP denotes a collection of information (objects) related to management domain in management information base (MIB). Multiple instances of the MIB objects are used by different network elements that are identified using SNMP context name and context engine id.</p>
V3 Authentication	<p>This field appears if you selected v3 in SNMP Version field.</p> <p>Select an authentication protocol from the drop-down list.</p> <p>Select <i>None</i> from the drop-down menu if you want to set SNMPv3 authentication to None.</p>
V3 Privacy	<p>This field appears if you selected v3 in SNMP Version field.</p> <p>Select a privacy protocol from the drop-down list.</p> <p>Select <i>None</i> from the drop-down menu if you want to set SNMPv3 privacy to None.</p>
V3 Auth Passphrase	<p>This field appears if you selected v3 in SNMP Version field.</p> <p>Enter a passphrase for SNMPv3 authentication.</p>
SNMPv3 Privacy Passphrase	<p>This field appears if you selected v3 in SNMP Version field</p> <p>Enter a passphrase to encrypt the ingest message.</p>

5. Click **OK** to save the configuration.

A confirmation window confirms that the edit operation was successful.

RELATED DOCUMENTATION

[About the Ingest Settings Page | 454](#)

Configure a Rule Using SNMP Scalar

The following example configuration of a rule has an SNMP tabular column as a scalar field and a second field with a scalar not indexed in the SNMP table.

To configure a rule with scalar fields:

1. Go to **Configuration > Rules**.

2. Click **+Add Rule** on top of the Rules page.

Fill in the topic and the rule name, description, and synopsis. You can see the new rule in the topic **external** by default, unless you specify a topic name here.

3. Click **+Add Sensor** and enter a name for the sensor.

4. Select **SNMP** as Sensor Type and *30s* as Frequency.

5. Click **Add Scalar** and enter **IF-MIB::ifNumber** in the scalar field.

You can also enter the OID of the scalar object.

6. Click **Add Scalar** and enter **IF-MIB::ifAdminStatus.16** in the scalar field.

NOTE: The index number is different from one device to another and from one system to another. Before you configure a rule for a device, get the index number of the scalar object in that device. For more information, see [SNMP index number](#).

7. Do one of the following:

- **Save** — Save your configuration changes but do not deploy the updated configuration. You can use this option when, for example, you are making several changes and want to deploy all your updates at the same time later.
- **Save and Deploy** — Save the rule configuration in the GUI and deploy the configuration on your production environment. The ingest starts collecting telemetry data based on the configuration changes.

RELATED DOCUMENTATION

[Configure SNMP Ingest | 499](#)

Configure SNMP Trap and Inform Notifications

IN THIS SECTION

- [Tasks You Can Perform | 504](#)
- [Find the Engine ID | 506](#)
- [Configure Trap Notifications | 506](#)
- [Configure Inform Notifications | 511](#)
- [Configure Port for Inform Notifications | 512](#)
- [Configure a Rule for SNMP Notification | 513](#)

Paragon Insights supports inform and trap notifications that devices send in the network for fault management. The *SNMP manager* (Paragon Insights) and the *SNMP agents* (devices) send traps and informs as notifications about change of state in network. Paragon Insights performs trigger evaluations on the traps and informs. Paragon Insights processes traps and informs from the configured device only if a playbook containing an SNMP-notification rule is running for the specified device. In all other cases, the SNMP Manager drops the trap or inform message.

The following sections describe relevant terms, configuration of traps and informs through CLI, port configuration, and accessing status of SNMP traps through CLI.

NOTE: You can configure SNMP trap notifications in SNMPv2c and SNMPv3. You can configure SNMP inform messages only when you use SNMPv3 protocol.

Tasks You Can Perform

Before you delve into SNMP trap and inform configurations, the following glossary can familiarize you with important concepts in SNMPv3 protocol.

The authoritative agent

In SNMPv3 transactions between two entities (agent and manager), Paragon Insights verifies the source device of notifications through authentication and

privacy. Authentication identifies and verifies the source of an SNMPv3 message. The privacy feature prevents packet analyzers from snooping the content of messages by encrypting the notification messages. The entity that controls the notification flow is known as an authoritative agent. In SNMPv3, the non-authoritative entity must know the *<Engine ID>* of the authoritative agent for a successful communication.

Traps or trap messages	A trap is an unacknowledged notification sent to the SNMP manager. In trap messages, SNMP agent is the authoritative agent. The administrator must configure the SNMP v3 <i><user></i> (distinct from the local IAM users) and <i><Context Engine ID></i> on the device that sends out the trap messages. For traps, the <i><Context Engine ID></i> is the <i>Engine ID</i> that uniquely identifies the SNMP agent.
Informs or inform messages	An inform is also a notification sent from an SNMP agent to the SNMP manager. In inform messages, SNMP manager is the authoritative agent. You configure the device that needs to send inform messages with the details of the remote authoritative agent, SNMP manager (Paragon Insights). The administrator must configure the <i><user></i> found in the remote SNMP manager.
Engine ID	<i><Engine ID></i> is a hexadecimal generated for a given agent that uniquely identifies the SNMP agent and needs to be unique across a given administrative domain. It also must be persistent across reboots or upgrades.
Security Engine ID	It is a security parameter in the SNMP communication between the agent and the manager. <i><Security Engine ID></i> is usually the <i><Engine ID></i> of the authoritative agent involved. A trap message has two parts: a header and a trap Protocol Data Unit (PDU). The header contains the <i><Security Engine ID></i> and a <i><username></i> set in the trap configuration. When an agent sends a trap, the parameters in the trap header are checked against the details in the USM table. The trap is further processed only when the parameters in the header matches with details in the USM table. In inform notifications, the <i><Security Engine ID></i> is Paragon Insight's <i>Engine ID</i> .
Context Engine ID	<i><Context Engine ID></i> is part of a trap PDU. It uniquely identifies a device which has sent the original trap message. <i><Context Engine ID></i> and <i><Security Engine ID></i> are identical in most cases.
USM Table	SNMP managers receiving the traps needs to maintain the USM table (User-based Security Model) which has <i><Security Engine ID></i> and <i><username></i> as the key to verify the source of the trap messages.

The following sections detail how to:

- ["Find the Engine ID" on page 506](#)
- ["Configure Trap Notifications" on page 506](#)

- ["Configure Inform Notifications" on page 511](#)
- ["Configure a Rule for SNMP Notification " on page 513](#)
- ["Configure Port for Inform Notifications" on page 512](#)

Find the Engine ID

Depending on if you configure devices to send trap or inform notifications, you need to first find the *<Engine ID>* of either the SNMP agent. You can refer to the sample commands below to find the engine id in Junos devices.

NOTE: The CLI command to find *<Engine ID>* varies from vendor-to-vendor.

To find the *<Engine ID>* of SNMP agents (devices) that are Junos-based platforms, enter the following command in CLI.

```
show snmp v3 engine-id
```

You will receive a HEX output as the device *<Engine ID>*.

Configure Trap Notifications

You can configure a device to send trap notifications using SNMPv2c and SNMPv3.

The source IP address needs to be unique across all the devices as it uniquely identifies the device. You can configure source IP address only for devices.

NOTE: In Paragon Insights, the SNMPv2c and SNMPv3 ingest and trap configurations share the same workflow.

To configure SNMP trap notifications at the device level:

1. Click the **Configuration > Device** option in the left navigation bar.
2. Select a device by clicking on the checkbox and click the edit device button (pencil icon).
The **Edit *Device-Name*** window appears.
3. Click on **Protocol > SNMP**.
4. Enter the necessary values in the text boxes and select the appropriate options for the device.

The following table describes the attributes in the **Edit *Device-Name*** window:

Attributes	Description
Protocols > SNMP	
Version	Select either v2c or v3 from the list.
Get Community (Only for SNMPv2c)	<p>Enter an SNMP Community string for SNMPv2c ingest.</p> <p>In SNMPv2c, the Community string is used to verify the authenticity of the ingest (request-response) message issued by the SNMP agent (devices such as routers, switches, servers, and so on).</p>
Port	Port number required for SNMP ingest (request-response) messages. The standard port number is 161 .
Timeout	<p>Enter the timeout period in seconds for SNMP notifications. You can enter a value between 0 and 65535.</p> <p>Timeout denotes the number of seconds after which the SNMP agent stops re-transmitting notification.</p>
Retry Count	<p>Enter a retry count between 0 and 255.</p> <p>Retry count is the number of times an SNMP agent attempt to retransmit notifications.</p>
V3 Username	Enter a username for SNMPv3 ingest (request-response), trap, and inform notifications.

(Continued)

Attributes	Description
V3 Context Name	<p>(Optional) Enter a context name for SNMPv3 trap and inform notifications.</p> <p>Context in SNMP denotes a collection of information (objects) related to management domain in management information base (MIB). Multiple instances of the MIB objects are used by different devices in the network. The devices are identified using the SNMP context name and the context engine ID.</p>
V3 Authentication	<p>This field appears if you selected v3 in SNMP Version field.</p> <p>Select an authentication protocol from the list.</p> <p>Select <i>None</i> from the list if you want to set SNMPv3 authentication to None.</p>
V3 Privacy	<p>This field appears if you selected v3 in SNMP Version field.</p> <p>Select a privacy protocol from the list.</p> <p>Select <i>None</i> from the list menu if you want to set SNMPv3 privacy to None.</p>
V3 Context Engine	<p>This field appears if you selected v3 in SNMP Version field.</p> <p>The <i>Engine ID</i> must be set to engine-id of the SNMP agent.</p>
V3 Auth Passphrase	<p>This field appears if you selected v3 in SNMP Version field.</p> <p>Enter a passphrase for SNMPv3 authentication.</p>

(Continued)

Attributes	Description
SNMPv3 Privacy Passphrase	<p>This field appears if you selected v3 in SNMP Version field</p> <p>Enter a passphrase to encrypt the ingest message.</p>
Device ID Details	
SNMP Source IPs	<p>Enter one or more IP address of source device. If a device has multiple IP addresses, separate them with a comma.</p> <p>The source IP address is used to identify the sender (SNMP agent) of trap and inform notifications.</p> <p>If you configured virtual IP address for SNMP trap receiver during installation, you can enter the virtual IP address here.</p>

5. Click **OK** to save the configuration.

A confirmation window confirms that the edit operation was successful.

In device groups, you can configure port number for traps and inform notifications. You can also configure log levels for SNMP notification.

1. Click the **Configuration > Device Groups** option in the left-navigation bar.
2. Select a device group and click on the edit button (pencil icon).

The Edit Device Group page appears.

3. Click **Advanced > Ports** to configure notification ports for traps and informs.
4. Click **Advanced > Logging > Service Logging Overrides** to configure SNMP logs.

The following table describes the attributes in the **Add a Device Group** window:

Table 82: Table 2: Add Device Group Page Details

Attributes	Description
Name	Name of the device group. (Required)
Description	Description for the device group.
Devices	<p>Add devices to the device group from the list. (Required)</p> <p>In Paragon Insights, you can add more than 50 devices per device group. However, the actual scale of the number of devices you can add depends on the available system resources.</p> <p>For example, let's say that you want to create a device group of 120 devices. In releases earlier than Release 4.0.0, we recommend that you create three device groups of 50, 50, and 20 devices respectively. With Paragon Insights, you just create one device group.</p>
Logging Configuration	
SNMP Notification	<p>Paragon Insights supports collecting log data for SNMP notification. You can collect different severity levels of logs for the snmp-notification service in a device group.</p> <p>Use these fields to configure which log levels to collect:</p> <p>Global Setting Log Level From the list, select the level of the log messages that you want to collect for every running Paragon Insights service for the device group. The level is set to None by default.</p> <p>Services Logging Overrides Select the log level from the list for any specific service that you want to configure differently from the Global Setting log level. The log level that you select for a specific service takes precedence over the Global Setting log configuration.</p>
Ports	
SNMP Notification Ports	<p>Enter port number(s) separated by comma, if you want to configure multiple ports. Paragon Insights listens on these ports for trap and inform notifications.</p>

5. Click **Save** to commit the configuration or click **Save and Deploy** to deploy the configuration in Paragon Insights.

Configure Inform Notifications

To enable devices to send inform notifications, you must configure SNMPv3 USM user(s).

To create USM users in Paragon Insights:

1. Go to **Configuration > Data Ingest > Settings**.
2. Select the **SNMP Notification** tab on the Ingest Settings page.
3. Click the **Usm Users** section.
4. Click the plus (+) icon to add a USM user.
5. In the Add USM User page, enter the username, and enable or disable the authentication and the privacy protocols by using the toggle button. .

If you disable Authentication and Privacy, the protocol and the passphrase fields do not appear.

NOTE: If you disable the Authentication protocol, the Privacy protocol cannot be enabled.

6. Click **Save** to only save the configuration, or click **Save and Deploy** to deploy the configuration in Insights.

After adding USM users, you can configure the following details in the Edit *Device-Name* page in Device Configuration and Edit Device Group page in Device Group Configuration.

Table 83: Table 3: SNMP Configuration for Informs in Device Groups

Attributes	Description
SNMP	
Version	You can set this field in Edit <i>Device-Name</i> page under Protocols > SNMP caret. Select v3 from the menu.
Port (Devices only)	You can set this field in Edit <i>Device-Name</i> page under Protocols > SNMP caret. Port number required for SNMP inform notifications. The standard port number for trap and inform notifications is 162 .

Table 83: Table 3: SNMP Configuration for Informs in Device Groups (Continued)

Attributes	Description
Notification Ports (Device Groups only)	<p>You can set this field in Edit Device Group page under Advanced > Ports > SNMP Notification Ports field.</p> <p>Enter notification ports separated by comma.</p> <p>Paragon Insights listens on the notification ports for traps and inform messages from device groups.</p>
Context Engine ID (Devices only)	<p>You can set this field in Edit Device-Name page under Protocols > SNMP caret.</p> <p>This field appears if you selected v3 in Version field.</p> <p>The <i>Engine ID</i> must be set to engine-id of the SNMP agent.</p>
Source IP Address (Devices only)	<p>You can set this field in Edit Device-Name page under Device Details ID > SNMP Source IPs caret.</p> <p>This field appears if you selected v3 in SNMP Version field.</p> <p>Enter the source IP address of the device. This field is optional.</p> <p>If you use NAT or an SNMP Proxy, the virtual IP address you configure for the SNMP Proxy must be set as the source IP address.</p>

Configure Port for Inform Notifications

By default, Paragon Insights listens for traps and informs in the standard SNMP trap port 162. If needed, you can change this port either at the global level (which is applicable to all device groups) or at the device group level applicable to a specific device group.

Port configured under ingest will apply to all device groups. Trap and Inform messages received through any other port are discarded.

To configure port number at the ingest level:

1. Go to **Configuration > Data Ingest > Settings** in the left-nav bar.
2. Select the **SNMP Notification** tab on the Ingest Settings page.
3. In the **Port** section, enter the port number.
4. Click **Save** to only save the configuration and **Save and Deploy** to deploy the configuration in Paragon Insights.

Port configured under device group will apply to only a specific device group. Traps and informs received through any other port are discarded. To configure port numbers at the device group level, see [Table 82 on page 510](#).

Configure a Rule for SNMP Notification

Once the device is configured to send traps or inform notification, you must configure a rule on the device with SNMP trap so that, Paragon Insights can process traps from the device. In device groups, you can apply a playbook instance that has the snmp-notification rule. When you configure SNMP notification in any rule, you must select the MIB name you want to monitor. Go to the [Juniper MIB Explorer](#) to browse MIB files for Junos OS devices and the [Cisco MIB Locator](#) to browse MIB files for Cisco devices.

The following example shows how you can configure a rule with SNMP notification to send alerts if an interface comes up for the *chassis.interfaces/* topic.

NOTE: It is assumed that you have configured the device or device group for SNMP trap notification. See "[Configure Trap Notifications](#)" on page 506 to configure SNMP trap notifications in devices or device groups.

To configure a rule under topic *chassis.interfaces/*:

1. Go to **Configuration > Rules**.
2. Click the **Add Rules** button in the Rules page.
Enter the rule name in the *topic/rule-name* format in the Rule field and description in the Description field. For example, *chassis.interfaces/linkup*.
3. Click **Add Sensor** button in Sensors tab.
4. Enter a name in the Sensor Name field and select **SNMP Notification** from the list in Sensor Type.
5. Enter notification name in *MIB-Name::Notification Name* format.

For example, **IF-MIB::linkDown**.

6. Click **Add Field** button in the Fields tab.

The fields for the SNMP Notification rule can be derived by any of the following methods:

- Variables (varbinds) for the given trap name.

The variables of the trap name can be defined as fields. The following steps use the example *IfAdminStatus* as varbind and *IF-MIB:linkDown* as the snmp-notification.

- a. Enter *IfAdminStatus* in the Field Name.
- b. Select *Integer* as **Field Type**.

The **Field Type** you enter in the GUI must be same as the type defined in the MIB File.

- c. Select *Sensor* as **Ingest Type** (field source).

The Ingest Type (field source) must be set to sensor.

- d. Select the sensor name from the list under **Sensor**.

The sensor name is the name you entered for the snmp-notification sensor.

- e. Enter *IfAdminStatus* as sensor path.

The **Path** must be set to the variable (varbind) name defined in the MIB file.

To add a second field for *IfOperStatus* as variable (varbind) for a given snmp-notification, follow the steps described above but change the field name and the sensor path to *IfOperStatus*.

- Notification name as a field name.

The SNMP notification name itself can be defined as field name to know which SNMP notification type is received. Some examples of SNMP Notification types are coldStart, warmStart, authenticationFailed, linkUp, and linkDown.

The following steps use the example *IfAdminStatus* as varbind and *IF-MIB:linkDown* as the snmp-notification.

- a. Enter *desired-name* in the Field Name.

- b. Enter *string* as **Field Type**.

The field type must be set to *string*.

- c. Select *Sensor* as **Ingest Type** (field source).

The Ingest Type (field source) must be set to sensor.

- d. Select the sensor name from the list under **Sensor**.

The sensor name is the name you entered for the snmp-notification sensor.

- e. Enter *_notification_name* as sensor path.

The **Path** must be set to '*_notification_name*'. *_notification_name* is a special path defined in Paragon Insights to get the notification name from the sensor data.

7. Click **Save** to commit the rule or **Save & Deploy** to deploy the rule in Paragon Insights.

You can see the new topic name and rule in the list of existing rules.

You can also configure triggers or functions based on the fields you add. See how to create a rule in GUI as explained in [Paragon Insights Rules and Playbooks](#).

You must include this rule in a playbook and apply the playbook's instance in a device or a device group.

To check the new SNMP notifications sent by device groups, log into Paragon Insights server as a root user and type the following command.

```
/var/local/healthbot/healthbot cli --device-group healthbot -s influxdb
```

You can track new entries of SNMP trap notifications. The notifications are sent to the Paragon Insights server for the fields (for example, IfAdminStatus) you configured.

RELATED DOCUMENTATION

| [Configure SNMP Ingest | 499](#)

Configure Outbound SSH Port for iAgent

In Paragon Automation Platform, you can configure TCP port number for iAgent (NETCONF) outbound SSH connections in device group configurations. Each device group has a unique port number. You can also configure a single TCP port number for outbound SSH traffic for all device groups at the ingest level.

To configure iAgent (NETCONF) port for all device groups:

1. Go to **Configuration > Data Ingest > Settings**.
2. Click the **SSH** tab on the Ingest Settings page.

You can use the toggle button to enable or disable the **Port** field.

3. Enter a port number for outbound SSH connections in the SSH page.

NOTE: If you configure the iAgent outbound SSH port number for a particular device group, the device group configuration takes precedence over the ingest configuration.

4. Do one of the following:
 - Click **Save and Deploy** to save and to deploy the configuration in your existing network configurations.
 - Click **Save** to only save the configuration but not deploy it in the existing network.

RELATED DOCUMENTATION

| [Add a Device Group](#) | 177

Configure System Log Ingest

IN THIS SECTION

- [Device Configuration](#) | 517
- [Configure Syslog Events Pattern](#) | 517
- [Add Patterns to a Pattern Set](#) | 520
- [Configure Header Pattern](#) | 520
- [Edit a Header Pattern](#) | 523
- [Clone a Syslog Events Pattern](#) | 523
- [Clone a Pattern Set](#) | 524
- [Configure Multiple Source IP Addresses for a Device](#) | 524

To be able to apply system log (syslog) ingest in a rule, you must first configure a device to send syslog data, configure syslog ingest by adding events pattern and applying patterns to pattern sets, and configure syslog header. You can refer this section to also clone pattern configurations and edit header configurations.

WHAT'S NEXT

When you configure syslog settings in Paragon Insights, you can opt to configure port, time zone for devices, and so on. For more information on optional configurations, see "[System Log Optional Configurations](#)" | 525.

Device Configuration

Configure your network device(s) to send syslog data to Paragon Insights. The example device configuration from the previous section is repeated here:

```
set system syslog host 10.10.10.1 any any
set system syslog host 10.10.10.1 allow-duplicates
set system syslog host 10.10.10.1 structured-data
```

10.10.10.1 = Load balancer IP address

Configure Syslog Events Pattern

An events pattern is a configuration to monitor some syslog event; you create a pattern for each event you want to monitor. This example uses patterns to monitor four syslog events (two structured and two unstructured).

NOTE: See the usage notes at the end of this section for more detail on what has been configured.

To configure Syslog pattern in GUI:

1. Click **Configuration > Data Ingest > Settings** in the left-nav bar.
2. Select the **Syslog** tab on the Ingest Settings page.
3. On the Syslog Settings page, click the plus (+) icon.
4. In the pop-up window that appears, enter the following parameters for the first pattern, named **snmp-if-link-down**:

ADD PATTERN ✕

Name * ✔

Event Id *

Description

Filter

Field + ✎ 🗑️

Name*	Description	From*	Type*	Key Field
<input type="radio"/> if-name		interface-name	integer	No
<input checked="" type="radio"/> snmp-index	<input style="width: 80px;" type="text"/>	snmp-interface-index	integer ▼	<input type="checkbox"/> No

2 items ✓ ✕

Constant + ✎ 🗑️

Cancel
Save
Save And Deploy

5. Click **Save and Deploy**.

6. Click the plus (+) once more and enter the following parameters for the second pattern, named **fpc-offline**:

ADD PATTERN ✕

Name * ✔

Event Id *

Description

Filter

Field + ✎ 🗑️

Name*	Description	From*	Type*	Key Field

NOTE: The full value entered in the Filter field is **fpc%(NUMBER:fpc) Marking ports % {WORD:port-status}**

7. Click **Save & Deploy**. On the Syslog Settings page you should see the two patterns you just created.

Usage notes for the patterns

For structured syslog:

- The event ID (**SNMP_TRAP_LINK_DOWN**) references the event name found within the syslog messages.
- Fields are optional for structured syslog messages; if you don't configure fields, the attribute names from the message will be treated as field names.
 - In this example, however we have user-defined fields:
 - The field names (**if-name**, **snmp-index**) are user-defined.
 - The field **interface-name** value is an attribute from the syslog message, for example, `ge-0/3/1.0`; this field is renamed as **if-name**
 - The field **snmp-interface-index** value is an attribute from the syslog message, for example, `ifIndex 539`; this field is renamed as **snmp-index**.
 - The field **snmp-interface-index** here is defined as an integer; by default the fields extracted from a syslog message are of type string, however type integer changes this to treat the value as an integer
- The constant section is optional, in this example, we have user-defined constants.
 - The constant name **ifOperStatus** is user-defined; in this case it has the integer value of '2'
- Filter configuration is optional for a structured syslog, though you can do so if desired; if used, the filter-generated fields will override the fields included in the syslog message.
- The key fields section is optional; by default the hostname and event ID will be the keys used by Paragon Insights; add additional key fields here; in this example, we have **key-fields**, namely **interface-name**, where the name and value are extracted from the syslog message's attribute-value pair

For unstructured syslog:

- The event ID is user defined. In this case, it is **PSEUDO_FPC_DOWN**
 - For example, neither the unstructured syslog `Nov 22 02:27:05 R1 fpc1 Marking ports down` nor its structured counterpart `<166>1 2019-11-22T02:38:23.132-08:00 R1 - - - fpc1 Marking ports down` includes an event ID.
- A filter must be used to derive fields (unlike proper structured syslog); this example uses `fpc%{NUMBER:fpc} Marking ports %{WORD:port-status}`, where **fpc** becomes the field name and **NUMBER** denotes the syntax used to extract the characters out of that particular portion of the message, for example "2".
 - An example of a syslog message that matches the grok filter is "fpc2 Marking ports down".
- Constant **fpc-status** - has a string value of 'online'.

Regarding filters:

- By default in a pattern, field and constant values are a string; to treat it as an integer or float, define the pattern's field type as integer or float.
- For unstructured patterns, you must configure a filter as the messages are sent essentially as plain text and don't include field info on their own.
- Filters should always be written to match the portion of message after the event ID; this allows the filter to parse a syslog message irrespective of whether it arrives in unstructured or structured format.
 - For example, the filter `fpc%{NUMBER:fpc} Marking ports %{WORD:port-status}` matches both versions of the following syslog message:
 - Structured: `<166>1 2019-11-22T02:38:23.132-08:00 R1 - - - fpc1 Marking ports down`
 - Unstructured: `Nov 22 02:27:05 R1 fpc1 Marking ports down`

Add Patterns to a Pattern Set

With the patterns configured, group them into a pattern set.

1. On the Syslog Settings page, click to expand the **Pattern Set** section and click the plus (+) icon.
2. In the pop-up window that appears, enter the following parameters:

3. Click **Save & Deploy**. On the Syslog Settings page you should see the pattern set you just created.

Configure Header Pattern

In Paragon Insights, you can configure the pattern for parsing the header portion of a syslog message. With this release, unstructured syslog messages of non-Juniper devices are supported. In earlier releases, you can only parse the payload portion of either a structured syslog message as specified in [RFC 5424](#) standard, or a Juniper device's unstructured syslog message.

In general, it is assumed that any unstructured syslog message matches the Juniper syslog message pattern. For example, you do not have to configure a Juniper header pattern as this pattern is inbuilt with Paragon Insights. However, in case of a non-Juniper device's unstructured syslog message that does not match with the inbuilt pattern, a first match is made with one of the user-configured header patterns. Following a successful match, the fields are extracted. When there is no match, the incoming syslog message is dropped.

To configure a header pattern:

1. Navigate to **Configuration > Data Ingest > Settings** in the left navigation bar.

The Ingest Settings page is displayed.

2. Click **Syslog** to view the Syslog Settings page.

3. Click to expand the **Header Pattern** section.

The **Header Pattern** section of the Syslog Setting page is displayed. You can add a new header pattern and edit or delete an existing header pattern from this page.

4. Click the plus (+) icon to add a new header.

The Add Header Pattern page is displayed.

5. Enter the following information in the Add Header Pattern page.

- a. Enter a name for the header pattern in the Name field.
- b. Enter a description for the header pattern in the Description field.
For example, you can provide a one-line description of why you are creating this header pattern.
- c. Enter the filter or regular expression (regex) for the header patter in the Filter field.

NOTE: You can use regex101.com to edit, validate, and modify the filter pattern you want to add to the header pattern.

An example of a filter pattern is `(.*):([A-Z][a-z]{2} \d{1,2} \d{1,2}:\d{1,2}:\d{1,2}\.d*)\s:\s([a-z]*)\[(\d*)\]:\s*(.*)\s*`.

- d. **log-host**, **log-timestamp**, and **log-message-payload** of the Fields section are mandatory fields that determine the position of the header.

In the Fields section,

- i. Click **log-host**, and enter the following information.
 - Enter a name for the log host in the Name field.
log-host is the default name.
 - Enter a description for log-host in the Description field.

The default description is **Position of host name**.

- Enter the capture group value with prefix **\$** in the From field.

The capture group determines from which position in the header the **log-host** starts.

ii. Click **log-timestamp**, and enter the following information.

- Enter a name for the log timestamp in the Name field.

log-timestamp is the default name.

- Enter a description for log-timestamp in the Description field.

The default description is **Position of time stamp**.

- Enter the capture group value with prefix **\$** in the From field.

The capture group determines from which position in the header the **log-timestamp** starts.

NOTE: Ensure that timestamp format follows this sample timestamp format: "Jan _2 15:04:05 2006". Otherwise parsing of syslog messages will lead to an undefined behaviour.

iii. Click **log-message-payload**, and enter the following information.

- Enter a name for the log message payload in the Name field.

log-message-payload is the default name.

- Enter a description for log-message-payload in the Description field.

The default description is **Position of payload**.

- Enter the capture group value with prefix **\$** in the From field.

The capture group determines from which position in the header the **log-message-payload** starts.

iv. (Optional) Click the plus (+) icon to add a new field.

- Enter a name for the new field in the Name field.
- Enter a description for the new field in the Description field.
- Enter the capture group value with prefix **\$** in the From field.

The capture group determines from which position in the header the new field starts.

You can add one or more than one fields by clicking plus (+) icon.

6. Enter the name(s) of key fields in the Key Fields field.
7. Click **Save** to save configuration, or click **Save & Deploy** to save and immediately deploy the configuration.

Alternatively, to cancel the configuration, click **Cancel**.

Edit a Header Pattern

To edit an already configured header pattern:

1. Navigate to **Configuration > Data Ingest > Settings** in the left-nav bar.
The Ingest Settings page is displayed.
2. Click **Syslog** to view the Syslog Settings page.
3. Click to expand the **Header Pattern** section.
The **Header Pattern** section of the Syslog Setting page is displayed.
4. Select the header pattern you want to edit by selecting the check box next to the name of the header pattern, and click the **Edit (pencil)** icon.
The Edit Header Pattern <Header Name> page is displayed.
5. After you have edited the required fields, click **Save** to save configuration.
You can also click **Save & Deploy** to save and immediately deploy the edited configuration.

Clone a Syslog Events Pattern

To clone an existing Syslog pattern:

1. Click **Configuration > Data Ingest > Settings** in the left-nav bar.
The Ingest Settings page is displayed.
2. Click the **Syslog** tab to view the Syslog Settings page.
3. Click to expand the **Events Pattern** section in the Syslog Settings page to view existing syslog events patterns.
4. To clone a pattern, click the **Clone**.
The Clone Pattern: <name of syslog pattern> page is displayed.

From the Clone Pattern: <name of syslog pattern> page, you can
 - Edit existing fields
 - Add new fields or constants
 - Add or remove key fields
5. Click **Save** to save configuration and clone the syslog pattern.
Alternatively, click **Save & Deploy** to save configuration, clone syslog pattern, and deploy the pattern.

Clone a Pattern Set

To clone an existing Syslog pattern set:

1. Click **Configuration > Data Ingest > Settings** in the left-nav bar.
The Ingest Settings page is displayed.
2. Click the **Syslog** tab to view the Syslog Settings page.
3. Click to expand the **PatternSet** section in the Syslog Settings page to view existing syslog pattern sets.
4. To clone a pattern set, click the **Clone**.

The Clone Pattern-set: *<name of pattern-set>* page is displayed.

From the Clone Pattern-set: *<name of pattern-set>* page, you can

- Edit the name and description fields
 - Add or remove patterns from the Patterns field.
5. Click **Save** to save configuration and clone the syslog pattern set.
Alternatively, click **Save & Deploy** to save configuration, clone syslog pattern set, and deploy the pattern set.

Configure Multiple Source IP Addresses for a Device

You can add additional source IP addresses for devices that send syslog messages using a different source IP address than the one originally configured in the Paragon Insights GUI.

To support additional source IP addresses:

1. Go to the **Configuration > Devices** page and click on the name of a device.
2. Click the **Edit (pencil)** icon.
The **Edit *Device-Name*** page appears.
3. Click **Device ID Details** and enter the IP address(es) in the **Syslog Source IPs** field.

RELATED DOCUMENTATION

| [Sensors Overview](#) | 429

System Log Optional Configurations

IN THIS SECTION

- [Configure Syslog Ports | 525](#)
- [Configure Syslog Time Zone | 525](#)
- [Configure Host Name Aliases for a Device | 526](#)

Configure Syslog Ports

By default, Paragon Insights listens for system log (syslog) messages from all device groups on UDP port 514. You can change the system-level syslog port, as well as configure one or more ports per device group. The more specific device group setting takes precedence over the system level setting.

To change the system-level syslog port:

1. Click **Configuration > Data Ingest > Settings** in the left navigation bar.
2. Select the **Syslog** tab on the left side of the page.
3. On the Syslog Settings page, edit the port number.
4. Click **Save & Deploy**.

To configure a syslog port for a device group:

1. Go to the **Configuration > Device Group** page and click on the name of a device group.
2. Click the **Edit (Pencil)** icon.
3. In the Edit Device Group window, click **Advanced > Ports** caret.
4. Enter the port(s) in the **Syslog Ports** field.
5. Click **Save & Deploy**.

Configure Syslog Time Zone

When a device exports structured syslog messages, time zone information is included within the message. However, unstructured syslog messages do not include time zone information. By default, Paragon Insights uses GMT as the time zone for a device. In these cases, you can assign a time zone to a device or device group within Paragon Insights.

To configure a device's time zone at the device level:

1. Go to the **Configuration > Devices** and click on the name of a device in the Device page.
2. Click the **Edit (Pencil)** icon.
The **Edit *Device-Name*** page appears.
3. Click **Advanced** drop-down menu and enter a value in the **Syslog Time Zone** field in the format **+/- hh:mm**. For example, **-05:00**.

Configure Host Name Aliases for a Device

When a device has more than one host name, such as a device with dual REs, syslog messages can arrive at the Paragon Insights server with a host name that is not the device's main host name. In these cases, you can add host name aliases for that device.

NOTE: If you add a device in Paragon Insights using its IP address, you must also add the host name that will appear in the syslog messages.

To configure additional hostname aliases:

1. Go to the **Configuration > Device** page and click on the name of a device.
2. Click the **Edit (Pencil)** icon.
The **Edit *Device-Name*** page appears.
3. Click **Device ID Details** and enter the host name(s) in the **Syslog Host Names** field.

RELATED DOCUMENTATION

[Configure System Log Ingest](#) | 516

Configure a Rule Using Syslog

With the syslog ingest settings complete, you can now create a rule using syslog as the sensor.

This rule includes three elements:

- A syslog sensor
- Four fields capturing data of interest
- A trigger that indicates when the interface goes down

NOTE: See the usage notes at the end of this section for more detail on what has been configured.

1. Click **Configuration > Rules** in the left-navigation bar.
2. On the Rules page, click the **+ Add Rule** button.
3. On the page that appears, in the top row of the rule window, set the rule name. In this example, it is **check-interface-status**.
4. Add a description and synopsis if you wish.
5. Click the **+ Add sensor** button and enter the following parameters in the Sensors tab:

The screenshot shows the 'Sensors' configuration interface. At the top, there are tabs for 'Sensors', 'Fields', 'Vectors', 'Variables', 'Functions', 'Triggers', and 'Rule Properties'. The 'Sensors' tab is selected. On the left, there is a list of sensors with 'if-status-sensor' highlighted. On the right, the configuration for 'if-status-sensor' is shown. The fields are: 'Sensor Name' (if-status-sensor), 'Sensor type' (Syslog), 'Pattern set' (check-interface-status), and 'Maximum hold period' (Enter syslog Maximum hold period). There is a '+ Add sensor' button at the top left and a 'Delete if-status-sensor' button at the top right.

6. Now move to the Fields tab, click the **+ Add field** button, and enter the following parameters to configure the first field, named **event-id**:

The screenshot shows the 'Fields' configuration interface. At the top, there are tabs for 'Sensors', 'Fields', 'Vectors', 'Variables', 'Functions', 'Triggers', and 'Rule Properties'. The 'Fields' tab is selected. On the left, there is a list of fields with 'event-id' highlighted. On the right, the configuration for 'event-id' is shown. The fields are: 'Field name' (event-id), 'Description' (Add a description for this field), 'Field type' (Field type), 'Add to rule key' (unchecked), 'Ingest type (Field source)' (Sensor), 'Sensor' (if-status-sensor), 'Path' (event-id), 'Zero suppression' (unchecked), and 'Data if missing' (Default value). There is a '+ Add field' button at the top left and a 'Delete event-id' button at the top right.

7. Click the **+ Add field** button again and enter the following parameters to configure the second field, named **fpc-slot**:

Sensors **Fields** Vectors Variables Functions Triggers Rule Properties

+ Add field **Delete fpc-slot**

event-id
fpc-slot
if-name
snmp-index

Field name* ?
fpc-slot

Description ?
Add a description for this field

Field type
Field type

Add to rule key ?

Ingest type (Field source)
Sensor

Sensor Path* ? Data if missing
if-status-sensor fpc Zero suppression Default value

8. Click the **+ Add field** button again and enter the following parameters to configure the third field, named **if-name**:

Sensors **Fields** Vectors Variables Functions Triggers Rule Properties

+ Add field **Delete if-name**

event-id
fpc-slot
if-name
snmp-index

Field name* ?
if-name

Description ?
Add a description for this field

Field type
Field type

Add to rule key ?

Ingest type (Field source)
Sensor

Sensor Path* ? Data if missing
if-status-sensor if-name Zero suppression all interfaces

9. Click the **+ Add field** button once more and enter the following parameters to configure the fourth field, named **snmp-index**:

The screenshot shows the 'Fields' configuration page. At the top, there are tabs for 'Sensors', 'Fields', 'Vectors', 'Variables', 'Functions', 'Triggers', and 'Rule Properties'. The 'Fields' tab is active. On the left, there is a sidebar with a '+ Add field' button and a list of field names: 'event-id', 'fpc-slot', 'if-name', and 'snmp-index'. The 'snmp-index' field is selected. The main form for the 'snmp-index' field is visible, with a 'Delete snmp-index' button in the top right. The form fields are: 'Field name*' (snmp-index), 'Description' (Add a description for this field), 'Field type' (Field type), 'Add to rule key' (toggle), 'Ingest type (Field source)' (Sensor), 'Sensor' (if-status-sensor), 'Path*' (snmp-index), 'Zero suppression' (toggle), and 'Data if missing' (Default value).

10. Now move to the Triggers tab, click the **+ Add trigger** button, and enter the following parameters to configure a trigger named **link-down**:

Sensors Fields Vectors Variables Functions **Triggers** Rule Properties

+ ADD TRIGGER

link-down DELETE LINK-DOWN

Trigger Name [?]
link-down

Frequency [?]
2s

Disable alarm deduplication

Term is-link-down

WHEN

Left operand	Operator	Right operand
\$event-id	=~	SNMP_TRAP_LINK-DOWN

Ignore Case

All in time range
300s

+ ADD CONDITION

THEN

Color
[Red]

Message
Link down for \$if-name(\$snmp-id: \$snmp-index)

Evaluate next term

Functions can be used as Trigger actions too, define them using the 'Functions' menu at the top.

Term is-fpc-down

WHEN

Left operand	Operator	Right operand
\$event-id	=~	PSEUDO_FPC_DOWN

Ignore Case

All in time range
300s

+ ADD CONDITION

THEN

Color
[Red]

Message
Link down for \$if-name of FPC\$ifc-slot

11. At the upper right of the window, click the **+ Save & Deploy** button.

Usage Notes for the rule

- Sensor tab

- The sensor name **if-status-sensor** is user-defined.
- The sensor type is **syslog**.
- Pattern set **check-interface-status** - it is assumed that the pattern set is configured earlier.
- If not set, the Maximum hold period defaults to 1s.
- Fields tab
 - Four fields are defined; although the patterns are capturing more than four fields of data, this example defines four fields of interest here; these fields are used in the trigger settings.
 - The field names (**event-id**, **fpc-slot**, **if-name**, **snmp-index**) are user-defined.
 - path **event-id** - default field created by syslog ingest in the raw table; references the field from the pattern configuration.
 - path **fpc** - references the value from the filter used in the unstructured pattern configuration.
 - path **if-name** - refers to the interface name field from the pattern configuration. See "[Configure System Log Ingest](#)" on page 516.
 - Data if missing **all interfaces** - if the if-name value is not included in the syslog message, use the string value "all interfaces".
 - path **snmp-index** - references the field from the pattern configuration.
- Triggers tab
 - The trigger name **link-down** is user-defined.
 - **frequency 2s** - Paragon Insights checks for link-down syslog messages every 2 seconds
 - term **is-link-down** - when **\$event-id** is like **SNMP_TRAP_LINK_DOWN**, in any syslog message in the last **300 seconds**, make **red** and show the message **Link down for \$if-name(snmp-id: \$snmp-index)**.
 - **\$event-id** - \$ indicates to reference the rule field **event-id**.
 - **Link down for \$if-name(snmp-id: \$snmp-index)** - for example, "Link down for ge-2/0/0 of FPC 2".
 - **\$if-name** - references the field value, i.e., the name of the interface in the syslog message.
 - term **is-fpc-down** - when **\$event-id** is like **PSEUDO_FPC_DOWN**, in any syslog message in the last 300 seconds, make **red** and show the message **Link down for \$if-name of FPC\$fpc-slot**.
 - **\$event-id** - \$ indicates to reference the rule field **event-id**.

- `$if-name` - “all interfaces”.
- Link down for `$if-name of FPC$fpc-slot` - for example, “Link down for all interfaces of FPC 2”.

RELATED DOCUMENTATION

[Configure System Log Ingest](#) | 516

[System Log Optional Configurations](#) | 525

Understand Inband Flow Analyzer 2.0

IN THIS SECTION

- [Benefits of Inband Telemetry Solution](#) | 532
- [Device Configuration](#) | 533
- [Paragon Automation Configuration](#) | 534

Inband Network Telemetry (INT) is a vendor-neutral network monitoring framework that provides per-hop granular data in the forwarding (data) plane. INT allows you to observe changes in flow patterns caused by microbursts, packet transmission delay, latency per node, and new ports in flow paths.

Inband Flow Analyzer (IFA) 2.0 is an implementation of INT in Junos OS switches. IFA enables the devices to collect flow data and export the data to external collectors for per-hop or end-to-end analysis. IFA uses probe packets to collect data such as per-hop latency, per-hop ingress and egress ports, packet Receive (RX) timestamp (in seconds), queue ID, congestion, and egress port speed. The IFA packets traverse the same path in the network and use the same queues as the packets in the forwarding plane. So, the IFA packets experience similar latency and congestion as the packets.

Benefits of Inband Telemetry Solution


- Samples flow data and exports data to collectors faster than traditional telemetry ingests.
- Gives a granular view of the source of fault, latency, and congestion in your live network.

Device Configuration

The QFX5120-32C and QFX5120-48Y devices support Inband Network Telemetry (INT) using IFA 2.0. The IFA probe packets collect flow metrics and export the data in the Internet Protocol Flow Information Export (IPFIX) format. Paragon Automation supports analysis of the IPv4 Virtual Extensible LAN (VXLAN) flow data using the IFA sensor. Paragon Automation identifies VXLAN flows if the standard VXLAN port 4789 is present as the destination port in the Outer L4 Header (Layer 4 Header). The format of the IFA 2.0 packet with the VXLAN flow data is shown in [Figure 42 on page 533](#).

NOTE: IFA uses revenue ports to export data to collectors. You cannot use management ports to export IFA data.

Figure 42: Format of VXLAN IFA 2.0 Packet



IFA probe packets use three nodes that have separate functionality as they collect flow information:

- *IFA Initiator Node (ingress node)*—Samples the IPv4 VXLAN traffic, converts packets to IFA format by adding an IFA header, and updates IFA probe packet with the Initiator Node metadata. The IFA Header has the total maximum length allowed for the IFA Metadata Stack. The metadata stack is where each node adds its respective hop-specific metadata.
- *IFA Transit Node*—Identifies IFA packets and appends metadata into the metadata stack of the packet. A transit node checks the current length against the total maximum length in the IFA Header.

If the current length equals or exceeds the maximum length, the Transit Node does not append its metadata to the IFA Metadata Stack.

- *IFA Terminating Node (egress node)*—Appends its metadata and exports a copy of the flow data to the IFA 2.0 application (the IFA firmware). The IFA application adds the egress port number, converts the packets into IPFIX format, and sends them to a collector such as Paragon Automation.

See [IFA 2.0 Probe for Real-Time Monitoring](#) for more information.

NOTE: You must configure the IFA Initiator Node, IFA Transit Node, and IFA Terminating Node in the QFX5120-32C and QFX5120-48Y switches.

Paragon Automation Configuration

In Paragon Automation, you must perform the following tasks:

1. Configure IP address of the deploy node and the UDP port in the device group. Paragon Insights deploys the IFA ingest on the configured deploy node.

See ["Add a Device Group" on page 177](#) for more information.

2. Configure one or more IFA flow IP addresses in devices. See ["Edit Devices" on page 168](#) for more information.

3. Create a rule for the IFA ingest.

See ["Configure a Custom Rule in Paragon Automation GUI" on page 343](#) for more information.

4. Create a playbook and deploy the playbook instance in device groups.

See ["Create a Playbook Using the Paragon Insights GUI" on page 309](#) to create a playbook in Paragon Automation.

See ["Manage Playbook Instances" on page 312](#) to deploy a playbook.

5. Configure device details such as device name and device ID in the ingest. See ["Configure Device Details for Inband Flow Analyzer Devices" on page 538](#).

Paragon Automation supports `hb_ifa_v2_0` as the IFA sensor name. The IFA sensor supports fields described in [Table 84 on page 535](#).

Table 84: IFA Sensor Fields

Field	Key Field	Data Type	Description
source_ip	Yes	String	IP address of the Initiator Node from which the IFA flow packets originate.
source_port	Yes	String	Source port of the Initiator Node from which the IFA packet originates.
dest_ip	Yes	String	IP address of the Terminating Node.
dest_port	Yes	String	Destination port of the Terminating Node that exports the IFA packets.
proto	Yes	String	Value of the protocol used for the IFA flow.
hop	Yes	String	<p>The hop field denotes the number of hops that the IFA packet traversed. If there are n nodes, the hop value starts with 1 for the Initiator node, 2 for the Transit node, and so on until it reaches the Terminating node that is assigned a value of n.</p> <p>NOTE: The IFA sensor can additionally assign the hop value 65,535 to describe end-to-end latency and the complete IFA flow path.</p> <p>In Paragon Automation rules, the hop field captures the sequence number (hop value) at each hop.</p>

Table 84: IFA Sensor Fields (Continued)

Field	Key Field	Data Type	Description
node_id	No	String	<p>Device ID of the IFA Initiator node, the IFA Transit node, or the IFA Terminator node, when the hop field's value is not 65,535. The device ID is present in the IFA Metadata Stack.</p> <p>When the hop field's value is 65,535, the node_id field denotes the complete path taken by the IFA probe packet.</p>
node_name	No	String	<p>Displays name of the IFA node associated with <i>node_id</i>, if you previously configured Paragon Automation to display the node_name.</p> <p>If you didn't configure Paragon Automation to display the node_name, the node_id is displayed.</p>
ingress_port	No	String	Ingress port of the node through which the IFA flow enters.
egress_port	No	String	Egress port of the node through which the IFA flow exits.
egress_portspeed	No	Unsigned integer 32	Speed (in Gigabits per second) of the egress port.

Table 84: IFA Sensor Fields (Continued)

Field	Key Field	Data Type	Description
congestion_bits	No	Unsigned integer 32	Congestion bit that indicates if an IFA packet experienced congestion or not.
queue_id	No	Unsigned integer 32	Identifier (ID) of the queue taken by the IFA packets in a node.
residence_time_ns	No	Unsigned integer 32	Time taken (in nanoseconds) by the IFA packet within a node.
rx_ts_ns	No	Unsigned integer 64	Receive time stamp value when the IFA probe packet enters a node.
latency	No	Unsigned integer 64	Difference between the received time stamp of the current node and the previous node, when the hop field's is not 65,535. When the hop field's value is 65,535, the latency field denotes the end-to-end latency of the complete path.

Paragon Automation ingests the IFA data as IPFIX records and creates multi-row entries in the time-series database (TSDB) for each IPFIX record. The TSDB rows capture per hop details such as:

- Ingress and egress ports
- Latency
- Receive packet (RX) time stamp
- Sequence number that increments at each hop
- A record of the end-to-end latency from the Initiator node to the Terminating node

RELATED DOCUMENTATION

[Add a Device Group | 177](#)

[Configure a Custom Rule in Paragon Automation GUI | 343](#)

Configure Device Details for Inband Flow Analyzer Devices

You can access the IFA Devices page from **Configuration>Data Ingest>Settings**. In the Ingest Settings page, click the **IFA Devices** menu.

In Paragon Automation, you can assign device specific details, such as device name and device ID, for the Inband Flow Analyzer (IFA) devices. Paragon Automation maps the device names to their respective device ID. The device ID you enter corresponds to the device ID in the Junos device configuration you earlier completed to enable IFA.

Monitoring the end-to-end path of IFA devices in a flow is more human readable. You can view the device name you configure in the IFA Devices page instead of the device ID. Paragon Automation displays this name as `node_name` in the time series database field table.

To assign a name to an IFA device:

1. Click the add icon (+).

The Create IFA Device page appears.

2. Enter the device ID.

Ensure that the device ID matches the device ID you entered for IFA devices. See [IFA 2.0 Probe for Real-Time Monitoring](#) for more information.

3. Do one of the following:

- a. Click **Save** to only save the configuration.

Paragon Automation saves the configuration to assign a name to the device but you cannot view the device name while monitoring the complete path of the IFA probe packets.

- b. Click **Save and Deploy** saves and deploys the configuration. You can see the device name you configured as `node_name` in the complete path of the IFA probe packets.

View the complete path of IFA probe packets in a flow using Grafana. See ["Grafana Overview" on page 423](#) for more information.

RELATED DOCUMENTATION

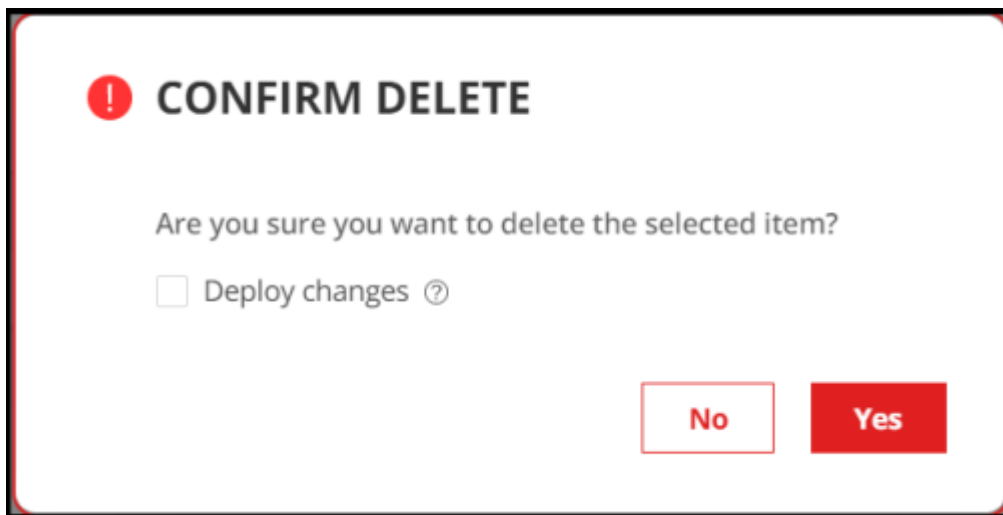
[Understand Inband Flow Analyzer 2.0 | 532](#)

Delete an Inband Flow Analyzer Device

To delete an Inband Flow Analyzer (IFA) Device:

1. Click **Configuration > Data Ingest > Settings** from the left-nav bar.
The Ingest Settings page is displayed.
2. Click the **IFA Device** tab to view the **IFA Devices** page.
3. Select the device that you want to delete, and click the **delete (trash can)** icon.
The **CONFIRM DELETE** pop-up appears.
4. Do one of the following:

Figure 43: Confirm Delete Pop-up



- Click **Yes** to delete the IFA device from the database. However, the changes are not applied to the ingest service.

NOTE:

- We recommended that you do not delete an IFA device that is currently in use.
- After you delete an IFA device from the database, you cannot associate that IFA device with another device group even if you have not deployed changes.

- You can also deploy changes to the ingest service or roll back the changes that you have already deleted, from the **Health Configuration Deployment Status** page. For more information, see "[Commit or Roll Back Configuration Changes in Paragon Insights](#)" on [page 185](#).

- Select the **Deploy changes** check box and then click **Yes** to delete the IFA device from the database, and to apply the changes to the ingest service.
- (Optional) Click **No** to cancel this operation.

The IFA device is deleted.

RELATED DOCUMENTATION

| [Understand Inband Flow Analyzer 2.0 | 532](#)

Understand Bring Your Own Ingest

IN THIS SECTION

- [Benefits | 541](#)

Paragon Automation provides Bring Your Own Ingest (BYOI) default plug-ins and support for BYOI custom plug-ins. BYOI plug-ins ingest telemetry data that is stored in third-party sources, such as a data lake or external databases. You can export such telemetry data that you collected through the BYOI plug-ins and store it in the Paragon Automation time series database (TSDB).

Bring Your Own Ingest Plugins include an input plugin that is developed by the user and an output plugin developed by Juniper Networks. The BYOI input plugin streams data from different data sources (Kafka) that use different data models (OpenConfig or NETCONF YANG), data encoding (based on Extended Markup Language [XML] or JavaScript Object Notation [JSON]), data security, and messaging buses (Kafka), and sends the data to the output plugin. The output plugin converts that data into the line protocol format and writes the data to the Paragon Automation TSDB.

Paragon Automation supports two types of plugins:

- *Default Plugin*—You can use default plugins to measure metrics that are unique to your network. You can work with Juniper Networks to develop default BYOI plug-ins. Juniper Network develops and sends you the default plug-ins, with the Kubernetes YAML files and the plugin configurations that are included as a compressed tar file. For more information on the workflow to deploy a default plug-in "[Bring Your Own Ingest Default Plug-in Workflow](#)" on page 127.
- *Custom Plugin*—You can use custom plug-in when you want to stream pre-existing telemetry data to Paragon Automation for analysis. You must develop the BYOI plug-in, build the BYOI plug-in ingest image, and load the plug-in image and Kubernetes YAML file for the plug-in to the Paragon Automation server.

Benefits

Deploying bring your own ingest plug-ins has the following benefits:

- Reduces the cost of collecting telemetry data from devices by reusing previously collected data to Paragon Automation.
- Enables you to use all Paragon Automation features—such as custom or default rules, playbooks, reports, graphs, network health view, and more—on the external data you ingest into Paragon Automation.

RELATED DOCUMENTATION

[Load BYOI Default Plug-ins | 541](#)

[Configure Bring Your Own Ingest Default Plug-in Instances | 542](#)

[Build and Load BYOI Custom Plug-in Images | 544](#)

Load BYOI Default Plug-ins

To load a default bring your own ingest (BYOI) plug-in in Paragon Automation, we recommend that you have a multinode installation of Paragon Automation in proof of concept and production systems. You can enter the commands to load your BYOI default plug-in on your server that hosts the Paragon Automation application.

1. Log into the primary node of the Paragon Automation using your server credentials.
2. Enter `cd /var/local/healthbot`.
3. Enter `./healthbot list-plugins -d`.
This command lists the loaded default plug-ins.
4. Enter `sudo ./healthbot -v load-plugin -n plugin-name`

This command loads a BYOI plug-in. For example, `./healthbot -v load-plugin -n tlive_kafka_oc`.

An authentication prompt appears.

5. Enter the GUI login credentials of Paragon Automation.

A CLI message confirms that the plug-in is successfully loaded.

6. Login to Paragon Automation GUI and click **Configuration>Data Ingest>Settings>BYO Ingest Plugins** to verify that the plugin is visible in the GUI.

You can see the loaded default plug-in listed in the Default Plugins page.

After the plug-in is successfully loaded, you must create an instance of the default plug-in in the Paragon Automation GUI and map the instance to existing device groups.

RELATED DOCUMENTATION

[Configure Bring Your Own Ingest Default Plug-in Instances | 542](#)

[Configure Ingest Mapping for Default BYOI Plug-in Instances | 558](#)

Configure Bring Your Own Ingest Default Plug-in Instances

Configure a new instance of the default Bring Your Own Ingest (BYOI) plug-in you earlier loaded.

To add an instance for the default plug-in:

1. Click **Configuration>Data Ingest>Settings**.

The Ingest Settings page appears.

2. Click the **BYO Ingest Plugins** tab.

If you loaded the default plug-in, you can see its name as a tab in the Default Plugins page.

For example, you can see a tab named Tlive-Kafka-oc in the Default Plugins page.

3. Click **+ New Instance** in Tlive-Kafka-oc tab.

The Add a tlive-kafka-oc Instance page appears.

4. Enter details as described in [Table 85 on page 543](#).

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

5. Do one of the following:

- Click **Save**—Save your ingest mapping but do not deploy the instance. You can use this option when you are making several changes and want to deploy all your updates at the same time later.
- Click **Save and Deploy**—Save and deploy the configuration.

Paragon Automation creates an instance for the BYOI plugin and deploys the configuration. The plug-in can now send telemetry data from an external source to Paragon Automation.

Table 85: Fields on the Add a Default-Plugin-Name Instance Page

Field	Description
Name	<p>Enter the name of the default plug-in instance.</p> <p>You can enter a name that follows the regular expression <code>'^[a-zA-Z][a-zA-Z0-9_-]*\$'</code> with maximum 64 characters.</p> <p>For example, t1 as name for an instance of Tlive-kafka-oc plug-in.</p>
Brokers	<p>Enter one or more broker names separated by commas.</p> <p>The default Kafka broker is healthbot.</p>
(Optional) Topics	<p>Enter healthbot as the topic to which the plugin subscribes.</p>
Authentication	
SASL Username	<p>Enter the username for authentication.</p> <p>You can enter a name that follows the regular expression <code>'^[a-zA-Z][a-zA-Z0-9_-]*\$'</code> with maximum 64 characters.</p>
Password	<p>Enter a password that has 6 to 128 characters long.</p> <p>The password must contain a combination of uppercase characters, lowercase characters, numbers, and special characters. Paragon Automation supports all keyboard special characters such as comma, asterisk, ampersand and so on.</p>
TLS	
CA Profile Name	<p>Enter the name of the trusted certificate authority.</p>
Local Certificate Profile Name	<p>Enter the name of the self-signed certificate you created for Paragon Automation.</p>
Skip Certification Chain and Host Verification	<p>Toggle the switch on if you want to skip verifying the integrity of the CA certificate.</p> <p>NOTE: The connection between the ingest plug-in and the external data source is encrypted even if you skip verification of the certificate.</p>

After you complete configuring the BYOI plug-in instance, you must map the default plug-in instance to existing device groups.

RELATED DOCUMENTATION

[Configure Ingest Mapping for Default BYOI Plug-in Instances | 558](#)

Build and Load BYOI Custom Plug-in Images

SUMMARY

IN THIS SECTION

- [Create a Process File for the Plug-in Image | 545](#)
- [Create a Shell Script for Configuration Updates | 548](#)
- [Tag and Export the BYOI Custom Plug-in Image | 548](#)
- [Configure Kubernetes YAML File | 549](#)
- [Assign Virtual IP Address to Plug-in | 552](#)
- [Load the BYOI Custom Plug-in | 554](#)

To send data to Paragon Automation using Bring Your Own Ingest (BYOI) Custom Plug-ins, you must code the input plug-in and create a plug-in ingest image file with a shell script and a process file. Paragon Automation executes the shell script when configurations of the BYOI ingest and device group (mapped to the ingest) change. The ingest image also contains a Kubernetes YAML file for the ingest container. The Kubernetes YAML file contains configurations that enable the Kubernetes engine to start and to stop the service for a BYOI plug-in ingest.

The workflow to build and to load the BYOI custom plug-in image is as follows:

1. Create a process file to write data to the Paragon Automation database and a shell script for configuration updates.
 - See "[Create a Process File for the Plug-in Image](#)" on page 545 for an example Python script (process file) that parses attributes in JSON configuration file to send data to the Paragon Insights database.

- See ["Create a Shell Script for Configuration Updates"](#) on page 548 for an example shell script.
2. Tag the image file and export the image as a compressed tar file. See ["Tag and Export the BYOI Custom Plug-in Image"](#) on page 548 for commands to tag and export the image file.
 3. Modify the Kubernetes Jinja template to create a YAML file for the Kubernetes container pod. The container pod is where the BYOI ingest is deployed in Paragon Automation. See ["Configure Kubernetes YAML File"](#) on page 549 for a sample Kubernetes Jinja template file.
 4. (Optional) Assign a different IP address if you want the BYOI plug-in to be accessible for external applications. See ["Assign Virtual IP Address to Plug-in"](#) on page 552 for a Kubernetes Jinja template in which you can assign a custom virtual IP address for the BYOI plug-in.
 5. Load the compressed tar file and Kubernetes YAML file to the Paragon Automation primary node. See ["Load the BYOI Custom Plug-in"](#) on page 554 to load the BYOI plug-in image file and the Kubernetes YAML file.

Create a Process File for the Plug-in Image

IN THIS SECTION

- [Decode Device Password | 548](#)

You can create a process file, such as the following example Python file, and include it in the BYOI plug-in image. When you run the image in a Kubernetes container, Paragon Automation executes the process file. The following sample Python file uses the attributes that are described in [Table 86 on page 547](#) to send a key (a random integer between 0 and 9 in the example file) for the measurement (topic/rule/sensor_name/byoi) to the database.

```
import requests
import time
import random
import os
import json

# read tand_host, tand_port from env vars
tand_host = os.environ.get('TAND_HOST', 'localhost')
tand_port = os.environ.get('TAND_PORT', '3000')
```

```

# read device, plugin, rule related attributes from config json
with open('/etc/byoi/config.json', 'r') as f:
    config_json = json.load(f)
# input, device, sensor as lists. modify index as needed. Using 0 for all idxes
database = config_json['hbin']['inputs'][0]['plugin']['config']['device'][0]['healthbot-storage']
['database']
measurement = config_json['hbin']['inputs'][0]['plugin']['config']['device'][0]['sensor'][0]
['measurement']

# Construct post request and data
url = 'http://{}:{}/write?db={}'. \
    format(tand_host, tand_port, database)
data = '{} {} {}'
metric = 'key'

while True:
    fields = '{}={}'.format(metric, random.randint(0,9))
    timestamp = int(time.time()) * 1000000000
    x = requests.post(url, data=data.format(measurement, fields, timestamp))
    time.sleep(10)

```

In the Python example file, use the following URL format to send data to Paragon Automation.

```
url = 'http://{}:{}/write?db={}'. \format(tand_host, tand_port, database)
```

The value for the database attribute must follow the syntax *database-name:device-group-name:device-name*.

The line protocol (post body) contains a string in the following format.

```
data = '{} {} {}'.format(measurement, fields, timestamp)
```

When you create an instance of a custom BYOI plug-in, a JavaScript Object Notation (JSON) configuration file is attached as a volume in the Kubernetes container for the BYOI ingest instance. The JSON configuration file contains information such as the device, device group, sensor path, hostname, and port of the backend service to which ingest data is sent. You can go through the JSON configuration using `/etc/byoi/config.json` file for all the available attributes.

[Table 86 on page 547](#) lists several key attributes in the JSON configuration file.

Table 86: Key Attributes in the JSON Configuration File

Attributes	Description	How to Access the Attributes
tand_host	Host name of the back-end service to which the plug-in sends the ingest data.	Environment Variable <i>\$TAND_HOST</i>
tand_port	Port number of the back-end service to which the plug-in sends data.	Environment Variable <i>\$TAND_PORT</i>
database	Name of database that stores the ingest data. The value for this attribute differs for each Paragon Automation node.	config_json['hbin'] ['inputs']['plugin']['config'] ['device']['idx']['healthbot-storage'] ['database']
measurement	Measurement of database in line protocol. For example, topic/rule/sensor_name/byoi NOTE: The value of sensor_name differs from sensor to sensor. See the https://docs.influxdata.com/influxdb to learn more about measurements.	config_json['hbin']['inputs'] ['plugin'] ['config']['device']['idx']['sensor'] [sensor_idx]['measurement']
fields	Metric-value pairs, separated by comma without space. For example, cpu_usage=50,memory_utilization=12.	None
timestamp	Unix Epoch timestamp in nanoseconds.	None

Table 86: Key Attributes in the JSON Configuration File *(Continued)*

Attributes	Description	How to Access the Attributes
password	<p>Encoded password of the device that receives the streaming data.</p> <p>See "Decode Device Password" on page 548 to decode the device password.</p>	<pre>config_json['hbin'] ['inputs']['plugin']['config'] ['device'][idx]['authentication'] ['password']['password']</pre>

Decode Device Password

The JSON configuration file can contain encoded sensitive information such as the password of the device that streams data.

To decode the data, you can initiate a POST call using the API `api-server:9000/api/v2/junos-decode` inside the plug-in container, with the encoded data in the post body.

The following sample POST call decodes an encoded password present in the JSON configuration file.

```
curl -X POST -L api-server:9000/api/v2/junos-decode -H "Content-Type: application/json" -d '{"data": "$ABC123"}' -v
```

Create a Shell Script for Configuration Updates

When the BYOI ingest image configuration or the Paragon Automation device group configuration changes, the JSON configuration file is updated. When a change in configuration occurs, Paragon Automation signals BYOI about the configuration update by executing the shell script located at `/jfit_scripts/jfit_reconfigure.sh`. When you build the ingest plug-in image, you must name your shell script `jfit_reconfigure.sh` and copy the script to `/jfit_scripts/` folder.

In the shell script, you can send a SIGHUP signal to the main process or simply kill old processes and start new ones. The following example shell script sends a SIGHUP signal to the main plug-in process:

```
pid=`ps -ef | grep ".*main.py" | grep -v 'grep' | awk '{ print $1}'` && \
kill -s HUP $pid
```

Tag and Export the BYOI Custom Plug-in Image

After you build the custom plug-in, you must tag the plug-in image and export it as a tar file. You can tag the plug-in image in the `healthbot_plugin_name:your_version` format. You must export the plug-in image as a compressed tar file using the following command:

```
docker save tag -o healthbot_<plugin_name>.tar.gz
```

Configure Kubernetes YAML File

The Kubernetes Jinja template file has the basic configuration required to deploy Kubernetes resources such as containers for the for the BYOI ingest pod.

You can use the following sample Kubernetes Jinja template to create a YAML file. You must replace:

- Placeholders for commands and args, `<ADD_COMMAND>` and `<ADD_ARGUMENTS>`.

For example, replace `<ADD_COMMAND>` with `python3` and `<ADD_ARGUMENTS>` with the name of your Python file.

- `<PLUGIN_NAME_CAPITALIZED>` with your plug-in name in upper case letters.

You can add other properties to the 'containers' part, such as volumes or Kubernetes secrets, in the template. After you modify the sample Kubernetes Jinja template, change the name of the file to `healthbot_<plugin_name>.yaml.j2` and save it.

The following code is a sample Kubernetes Jinja template.

```
set sg_name = '-' + env['SUBGROUP'] -%}
{%- set sg_dir = '_' + env['SUBGROUP'] -%}
{% if env['SUBGROUP'] == '' -%}
  {%- set sg_name = '' -%}
  {%- set sg_dir = '' -%}
{%- endif %}
kind: ConfigMap
apiVersion: v1
metadata:
  namespace: {{ env['NAMESPACE'] }}
  name: {{ env['GROUP_TYPE'] }}-{{ env['GROUP_NAME_VALID'] }}
    {{ sg_name }}-{{ env['CUSTOM_PLUGIN_NAME'] }}
  labels:
    app: {{ env['CUSTOM_PLUGIN_NAME'] }}
    group-name: {{ env['GROUP_NAME'] }}
    group-type: {{ env['GROUP_TYPE'] }}
    subgroup: {{ env['SUBGROUP'] }}
data:
  TAND_HOST: '{{ env['GROUP_TYPE_SHORT'] }}-{{ env['GROUP_NAME_VALID'] }}
    {{ sg_name }}-{{ env['CUSTOM_PLUGIN_NAME'] }}-terminus'
  TAND_PORT: '{{env['tand:TAND_PORT']}}'
  PUBLISHD_HOST: '{{env['publishd:PUBLISHD_HOST']}}'
```

```

PUBLISHD_PORT: '{{env["publishd:PUBLISHD_PORT"]}}'
CONFIG_MANAGER_PORT: '{{env["configmanager:CONFIG_MANAGER_PORT"]}}'
CHANNEL: '{{ env["GROUP_TYPE"] }}-{{ env["GROUP_NAME"] }}'
GODEBUG: 'madvdontneed=1'
IAM_SERVER: '{{ env["iam:IAM_SERVER"] }}'
IAM_SERVER_PORT: '{{ env["iam:IAM_SERVER_PORT"] }}'
IAM_SERVER_PROTOCOL: '{{ env["iam:IAM_SERVER_PROTOCOL"] }}'
IAM_NAMESPACE: '{{ env["iam:IAM_NAMESPACE"] }}'

---
apiVersion: apps/v1
kind: Deployment
metadata:
  namespace: {{ env["NAMESPACE"] }}
  name: {{ env["GROUP_TYPE"] }}-{{ env["GROUP_NAME_VALID"] }}
    {{ sg_name }}-{{ env["CUSTOM_PLUGIN_NAME"] }}
  labels:
    app: {{ env["CUSTOM_PLUGIN_NAME"] }}
    group-name: {{ env["GROUP_NAME"] }}
    group-type: {{ env["GROUP_TYPE"] }}
    subgroup: {{ env["SUBGROUP"] }}
spec:
  replicas: 1
  selector:
    matchLabels:
      app: {{ env["CUSTOM_PLUGIN_NAME"] }}
      group-name: {{ env["GROUP_NAME"] }}
      group-type: {{ env["GROUP_TYPE"] }}
      subgroup: {{ env["SUBGROUP"] }}
  template:
    metadata:
      namespace: {{ env["NAMESPACE"] }}
    labels:
      app: {{ env["CUSTOM_PLUGIN_NAME"] }}
      group-name: {{ env["GROUP_NAME"] }}
      group-type: {{ env["GROUP_TYPE"] }}
      subgroup: {{ env["SUBGROUP"] }}
    spec:
      tolerations:
        - key: "node.kubernetes.io/not-ready"
          operator: "Exists"
          effect: "NoExecute"
          tolerationSeconds: 180

```

```

- key: "node.kubernetes.io/unreachable"
  operator: "Exists"
  effect: "NoExecute"
  tolerationSeconds: 180
initContainers:
- name: sync
  image: {{ env['REGISTRY'] }}/{{ env['HEALTHBOT_INIT_CONTAINER_IMAGE'] }}:
    {{ env['HEALTHBOT_INIT_CONTAINER_TAG'] }}
  imagePullPolicy: Always
  command: ["python3"]
  args: ["/root/sync_files.py", "-c", "{{ env['GROUP_TYPE'] }}-
    {{ env['GROUP_NAME'] }}" ]
  env:
  - name: NODE_IP
    valueFrom:
      fieldRef:
        fieldPath: status.hostIP
  envFrom:
  - configMapRef:
      name: {{ env['GROUP_TYPE'] }}-{{ env['GROUP_NAME_VALID'] }}
        {{ sg_name }}-{{ env['CUSTOM_PLUGIN_NAME'] }}
containers:
- name: {{ env['CUSTOM_PLUGIN_NAME'] }}
  image: {{ env['REGISTRY'] }}/{{ env['HEALTHBOT_&lt;variable
>PLUGIN_NAME_CAPITALIZED</variable>&gt;_IMAGE&rsquo;] }}:
    {{ env['&rsquo;HEALTHBOT_&lt;variable>PLUGIN_NAME_CAPITALIZED</variable
>&gt;_TAG&rsquo;] }}
  imagePullPolicy: Always
  #example
  #command: ["python3"]
  #args: ["/main.py"]
  command: [&lt;variable>ADD_COMMAND</variable>&gt;]
  args: [&lt;variable>ADD_ARGUMENTS</variable>&gt;]
  env:
  - name: NODE_IP
    valueFrom:
      fieldRef:
        fieldPath: status.hostIP
  envFrom:
  - configMapRef:
      name: {{ env['GROUP_TYPE'] }}-{{ env['GROUP_NAME_VALID'] }}
        {{ sg_name }}-{{ env['CUSTOM_PLUGIN_NAME'] }}
volumeMounts:

```

```

- name: default
  mountPath: /etc/byoi
- name: data-model
  mountPath: /etc/ml
volumes:
- name: default
  hostPath:
    type: DirectoryOrCreate
    path: {{ env['JFIT_OUTPUT_PATH'] }}/{{ env['GROUP_NAME'] }}
        {{ sg_dir }}/custom_{{ env['CUSTOM_PLUGIN_NAME'] }}_collector
- name: data-model
  hostPath:
    type: DirectoryOrCreate
    path: {{ env['JFIT_ETC_PATH'] }}/data/models/{{ env['GROUP_NAME'] }}
imagePullSecrets:
- name: registry-secret

```

Assign Virtual IP Address to Plug-in

For a custom BYOI plug-in to be reachable from an external network, the plug-in needs to be exposed as a Kubernetes loadbalancer service. This is an optional configuration. By default, the plug-in uses virtual IP address of the Paragon Automation gateway. You can also assign a custom virtual IP address and add the following template to the end of the Kubernetes Jinja template file in "[Configure Kubernetes YAML File](#)" on page 549.

Ensure that you replace `<PLUGIN_PORT>` and `<PROTOCOL>` in the given template with your desired values such as port 80 for protocol HTTP. See the <https://kubernetes.io/docs/concepts/services-networking/service/#protocol-support> for supported protocols.

To configure a custom virtual IP address for the BYOI plug-in, replace `<custom_load_balancer_ip>` with a virtual IP address.

```

---
{% set service_values = env.get('SERVICE_VALUES', {}) -%}
{%- set global_annotations = service_values.get('annotations') -%}
{%- set global_load_balancer_ip = service_values.get('loadBalancerIP') -%}
{%- set custom_annotations = service_values.get(svc_name, {}).get('annotations') -%}
{%- set custom_load_balancer_ip = service_values.get(svc_name, {}).get('loadBalancerIP') -%}
{%- set service_type = service_values.get(svc_name, {}).get('type', 'LoadBalancer') -%}
{%- for ip in env['LOAD_BALANCER_IPS'] %}
apiVersion: v1
kind: Service

```

```

metadata:
  namespace: {{ env['NAMESPACE'] }}
  {%- if loop.index0 == 0 %}
  name: {{ env['GROUP_TYPE_SHORT'] }}-{{ env['GROUP_NAME_VALID'] }}
        {{ sg_name }}-{{ env['CUSTOM_PLUGIN_NAME'] }}
  {%- else %}
  name: {{ env['GROUP_TYPE_SHORT'] }}-{{ env['GROUP_NAME_VALID'] }}
        {{ sg_name }}-{{ env['CUSTOM_PLUGIN_NAME'] }}-{{loop.index0}}
  {%- endif %}
  labels:
    app: {{ env['CUSTOM_PLUGIN_NAME'] }}
    group-name: {{ env['GROUP_NAME'] }}
    group-type: {{ env['GROUP_TYPE'] }}
    subgroup: '{{ env['SUBGROUP'] }}'
  annotations:
    {% if env.get('LOADBALANCER_PROVIDER', 'User') == 'HealthBot' %}
    metallb.universe.tf/allow-shared-ip: healthbot-{{loop.index0}}
    {% elif custom_annotations %}
    {{ custom_annotations }}
    {% elif global_annotations %}
    {{ global_annotations }}
    {% else %}
    {}
    {% endif %}
spec:
  type: LoadBalancer
  {%- if env.get('LOADBALANCER_PROVIDER', 'User') == 'HealthBot' %}
  loadBalancerIP: {{ ip }}
  {%- elif custom_load_balancer_ip %}
  loadBalancerIP: {{ <custom_load_balancer_ip> }}
  {%- elif global_load_balancer_ip %}
  loadBalancerIP: {{ global_load_balancer_ip }}
  {%- endif %}
  selector:
    app: {{ env['CUSTOM_PLUGIN_NAME'] }}
    group-name: {{ env['GROUP_NAME'] }}
    group-type: {{ env['GROUP_TYPE'] }}
    subgroup: '{{ env['SUBGROUP'] }}'
  ports:
  - name: port
    port: <PLUGIN_PORT>

```

```
protocol:<PROTOCOL>
{% endfor %}
```

After you configure the port and protocol, external applications can communicate with the custom BYOI plug-in via `<gateway_IP>:<PLUGIN_PORT>`. If you configured a custom virtual IP for the plug-in to act as a loadbalancer service, external applications can communicate with the plug-in via `<custom_load_balancer_ip>:<PLUGIN_PORT>`.

Use `<0.0.0.0>:<PLUGIN_PORT>` to connect to the server running inside the plug-in from the Kubernetes host.

You can configure different ports for different applications in the given Kubernetes Jinja template under the ports section.

NOTE: If you configure different port numbers in the Kubernetes Jinja template, you must hard code the port number and the corresponding port name in your respective applications.

Load the BYOI Custom Plug-in

Mount the BYOI custom plug-in image tar file and the modified Kubernetes YAML file to the Paragon Automation primary node and load the plug-in in the Paragon Automation management CLI.

1. Mount the BYOI plug-in image (compressed tar file) using the following command:

```
export HB_EXTRA_MOUNT1=/path/to/healthbot_plugin_name.tar.gz
```

2. Mount the Kubernetes YAML file using the following command:

```
export HB_EXTRA_MOUNT2=/path/to/healthbot_plugin_name.yml.j2
```

3. Load the plug-in image and the Kubernetes YAML file using the following command

```
sudo -E healthbot load-plugin -i $HB_EXTRA_MOUNT1 -c $HB_EXTRA_MOUNT2
```

You can see a confirmation message when the plug-in loads successfully.

4. (Optional) Select **Configuration > Data Ingest > Settings > BYO Ingest Plugins** page and view the custom plug-in in the **Custom Plugins** tab.

After you load your plug-in, create an instance of the custom BYOI plug-in in the Bring Your Own Ingest page. Since custom plug-ins do not use default Paragon Automation resources, you must configure a new rule and a playbook for the ingest plug-in.

SEE ALSO

[Configure Bring Your Own Ingest Custom Plug-in Instances](#) | 555

Configure Bring Your Own Ingest Custom Plug-in Instances

For custom ingest plug-ins, you must configure key-value pairs in the Paragon Automation GUI. To configure a custom ingest plug-in instance:

1. Click **Configurations > Data Ingest > Settings > BYO Ingest Plugins**.

The Bring Your Own Ingest Plugin page appears.

2. Click the add icon (+) in the Custom tab.

The Create Custom Plugins page appears.

3. Enter the details as described in [Table 87 on page 555](#).

4. Do one of the following:

- a. **Save**—Save your ingest mapping but do not deploy the updated configuration. You can use this option when, for example, you are making several changes and want to deploy all your updates at the same time later. See ["Commit or Roll Back Configuration Changes in Paragon Insights" on page 185](#) for more information.
- b. **Save & Deploy**—Save the custom plugin instance configuration and deploy the configuration in Paragon Automation.

Fields marked with an asterisk (*) are mandatory.

Table 87: Fields on the Create Custom Plugins Page

Fields	Description
Name	Enter the name of the custom ingest plug-in instance.
Plugin Name	Enter the name given to the input ingest plugin that you created.
Service Name	Enter a name for the service. The service name can be the same as the name of the custom plug-in.
Key	Name of the key parameters in the ingest plug-in. For example, sensor or frequency. You can add one or more ingest parameters as a key using the + icon.

Table 87: Fields on the Create Custom Plugins Page (*Continued*)

Fields	Description
Value	Enter the value for the key parameter. For example, you must enter a sensor path as value if sensor is your key.
Authentication	
SASL Username	Enter the username for authentication.
Password	Enter a password. The password must be 6 to 128 characters long and must contain a combination of uppercase and lowercase characters. It must also contain numbers and special characters.
TLS	
CA Profile Name	Enter the name of the trusted certificate authority (CA).
Local Certification Profile Name	Enter the name of the self-signed certificate you created for Paragon Automation.
Skip Certification Chain and Host Verification	Toggle the switch on if you want to skip verifying the integrity of the CA certificate. NOTE: The connection between the ingest plug-in and the external data source is encrypted, even if you skip verification of the certificate.

After you configure a custom plug-in instance, you must create a rule with the BYOI ingest parameters. To collect data using the custom ingest plug-in, you must add this rule to a new playbook and run the playbook's instance on device groups or network groups.

RELATED DOCUMENTATION

[Use Sample Rule and Playbook Configurations for BYOI Custom Plug-in Instances](#) | 557

Use Sample Rule and Playbook Configurations for BYOI Custom Plug-in Instances

Before you configure a custom rule, playbook, and device group for the Bring Your Own Ingest (BYOI) custom plug-in, you must create an instance of the custom plug-in. See ["Configure Bring Your Own Ingest Custom Plug-in Instances" on page 555](#) for more information.

You must write your own rules for BYOI custom plug-ins. You must deploy a custom rule in a new playbook that you create for the BYOI plug-in ingest. You must also create a device group that you add later in the BYOI playbook. Enter the rule and playbook configurations in Paragon Automation management CLI.

The following steps provide the rule and playbook configurations that you enter in the Rule Builder CLI in Paragon Insights.

1. Select **Configuration>Rules** in the Paragon Automation GUI.

The Rules page appears.

2. Click the **Rule Builder CLI** to open the management CLI.

The Rule Builder CLI terminal appears with the `user@host` configuration prompt. User denotes a root or a non-root user.

3. Type `configure` to enter the configuration mode.

You can see a message that confirms that you are in configuration mode.

4. Enter the following command to configure a new rule for the BYOI custom plug-in.

The following sample rule triggers an alert when the value of `key` is greater than 5. The trigger alert frequency is set to 10 seconds.

```
set healthbot topic external rule r1 sensor sensor_a byoi plugin name example-plug-in
set healthbot topic external rule r1 field key sensor sensor_a path key
set healthbot topic external rule r1 field key type integer
set healthbot topic external rule r1 trigger trigger_789 frequency 10s
set healthbot topic external rule r1 trigger trigger_789 term Term_1 when greater-than "$key"
5
set healthbot topic external rule r1 trigger trigger_789 term Term_1 then status color red
set healthbot topic external rule r1 trigger trigger_789 term Term_1 then status message BAD
set healthbot topic external rule r1 trigger trigger_789 disable-alarm-deduplication
```

NOTE:

- The *example-plug-in* in the sample rule refers to the plug-in name you entered in the **Name** field, when creating an instance of the custom plug-in.
- The *sensor_a* in the sample rule refers to the value configured for the key parameter in the **Value** field, when creating an instance of the custom plug-in.

5. Type `commit`.

You can see a message that confirms that the commit is complete.

6. Enter the following command to configure a new playbook.

Use the following sample command to create a playbook `p1` for the ingest that uses the sample rule `r1` that you previously created.

```
set healthbot playbook p1 rules external/r1
```

7. Type `commit`.

You can see a message that confirms that the commit is complete.

8. Type `exit` to exit the configuration mode.

9. In the operational mode, type `request healthbot deploy` to deploy the configuration.

The rules list on the Rules page refreshes automatically and shows the new rule you deployed. To see the new playbook, see the Playbooks page using the **Configuration > Playbooks** menu.

After you deploy the playbook, you can monitor the health of the BYOI ingest in the Network Health page (**Monitoring > Network Health**).

RELATED DOCUMENTATION

[Build and Load BYOI Custom Plug-in Images | 544](#)

Configure Ingest Mapping for Default BYOI Plug-in Instances

For default ingest plug-ins, you must configure the sensors and device groups that can use the plug-in.

To configure the ingest mappings for a default plug-in:

1. Go to **Configuration > Data Ingest > Settings > BYOI Plugins**.

The Ingest Settings page appears.

2. Click **+New Mapping** in the Ingest Mapping tab.

The Add a New Ingest Mapping page appears.

NOTE: You can configure multiple ingest mappings for the same default plug-in.

3. Enter the details as described in [Table 88 on page 559](#).
4. Do one of the following:
 - **Save**—Save your ingest mapping but do not deploy the updated configuration. You can use this option when, for example, you are making several changes and want to deploy all your updates at the same time later.
 - **Save & Deploy**—Save the ingest mapping configuration and deploy the configuration in your production environment.

You can use BYOI ingest in rules after you configure the ingest mapping for a default plug-in.

Table 88: Attributes in Add a New Ingest Mapping Page

Attributes	Description
Name	<p>Enter a name for the ingest mapping.</p> <p>You can enter a name that follows the regular expression <code>'^[a-zA-Z][a-zA-Z0-9_-]*\$'</code> with maximum 64 characters.</p> <p>The name is an instance identifier of the ingest mapping.</p>
Sensor Type	Select the type of sensor to be used for the plug-in from the list.
Plugin Name	Enter the name of the default ingest plug-in instance.
Constraint to Device Groups	Select device groups from the list. The ingest and sensor mapping is applied to only the selected device groups.

RELATED DOCUMENTATION

[Understand Bring Your Own Ingest | 540](#)

Delete a BYOI Plug-in

To delete a BYOI plug-in:

1. Click **Configuration > Data Ingest > Settings** from the left-nav bar.
The Ingest Settings page is displayed.
2. Click the **BYO Ingest Plugins** tab to view the **Bring Your Own Ingest** page.
3. Do one of the following.


NOTE: Default plug-ins cannot be deleted.

- To delete an ingest mapping:
 - a. Click **Ingest Mapping**.
 - b. Select the ingest mapping that you want to delete.
 - c. Click the **delete (trash can)** icon.
- To delete a custom plug-in:
 - a. Click **Custom Plugins**.
 - b. Select the custom plug-in that you want to delete.
 - c. Click the **delete (trash can)** icon.

The **CONFIRM DELETE** pop-up appears.

4. Do any one of the following:

Figure 44: Confirm Delete Pop-up



- Click **Yes** to delete BYOI plug-in information from the database. However, the changes are not applied to the ingest service.

NOTE:

- We recommended that you do not delete a BYOI plug-in that is currently in use.
- After you delete a BYOI plug-in from the database, you cannot map that BYOI plug-in with another device group even if you have not deployed changes.
- You can also deploy changes to the ingest service or roll back the changes that you have already deleted, from the **Health Configuration Deployment Status** page. For more information, see "[Commit or Roll Back Configuration Changes in Paragon Insights](#)" on [page 185](#).

- Select the **Deploy changes** check box and then click **Yes** to delete BYOI plug-in information from the database, and to apply the changes to the ingest service.
- (Optional) Click **No** to cancel this operation.

The BYOI plug-in information is deleted.

About the Diagnostics Page

Paragon Insights supports four verification and troubleshooting features:

- Paragon Insights Self Test – To know how to run the test, see ["Use the Self Test Tool"](#) on page 564.
- Device Reachability Test – To know how to run the test, see ["Use the Reachability Test"](#) on page 566.
- Ingest Connectivity Test – To know how to run the test, see ["Use the Ingest Test Tool"](#) on page 567.
- Debug No-Data – To know how to run the test, see ["Use the No-Data Tool"](#) on page 568

Paragon Insights Self Test

The self-test tool validates the core functionality of Paragon Insights. To perform the self test, the tool performs a typical set of tasks:

- Adds a simulated device to Paragon Insights
- Creates a device group, and adds the device
- Creates a rule
- Creates and deploys a playbook
- Streams data from the simulated device
- Displays ongoing status in the dashboard

The self-test instance essentially acts as a fully working setup, running entirely within the Paragon Insights system. When testing is complete, the tool provides a report. In addition to validating the Paragon Insights installation, the self-test feature also provides:

- An easy way to do a quick demo - the self test instance provides a simulated device connected to Paragon Insights, so you can demo Paragon Insights with no need to add a real device or apply playbooks.
- A good way for new users to get started - the self test auto-configures a simulated device connected to Paragon Insights, thereby eliminating the complexity of adding devices, applying playbooks, and so on.
- A 'running reference' - if there is an issue with real devices, you can use a self-test instance to help determine where the issue is; if the self-test instance is OK then the problem is not with the Paragon Insights system.

Device Reachability Test

Earlier, you would need to get through the entire setup procedure - add the device, setup a device group, apply playbooks, monitor devices - at which point the health pages would indicate "no data" indicating that the setup did not work correctly. But, "no data" does not indicate whether the problem is a reachability issue or data streaming issue.

In Paragon Insights, the device reachability tool can verify connectivity to a device. The tool performs tests using ping and SSH. Paragon Insights uses the device's IP address or host name, based on what was configured when adding the device.

Ingest Connectivity Test

The ingest connectivity tool can verify ingest methods where Paragon Insights initiates the connection, such as OpenConfig, iAgent, and SNMP. Paragon Insights does not test UDP-based ingest methods, such as syslog and Native GPB, as the UDP parameters are common to a device group and not specific to a device. This tool provides multiple benefits such as:

- Helps to identify when there might be missing configuration on the network device side.
- Helps you choose appropriate playbooks and rules that use sensors compatible with the supported ingest methods.
- Helps to identify ingest connectivity issues early on, rather than troubleshoot the “no-data” issue described in the previous section.

Paragon Insights validates each supported ingest method in its own way:

- **OpenConfig:** Establishes a gRPC connection with the device using its IP/host name, gRPC port, and credentials
- **iAgent:** Establishes a NETCONF session with the device using its IP/host name, NETCONF port, and credentials
- **SNMP:** Executes a simple SNMP GET command; expects to get a reply from the device

Debug No- Data

The debug no-data tool helps to determine why a device or rule is showing a status of “no-data”. The tool takes a sequential, step-by-step approach to determine at which stage incoming data is getting dropped or blocked, as follows:

- **Paragon Insights Services** – Verify that all common and device group-related services are up and running
- **Device Reachability** – Test connectivity to device using ping and SSH
- **Ingest Connectivity** – Verify that the configured ingest session is established
- **Raw Data Streaming** – Verify whether the ingest is receiving any raw data from the devices
- **Field Processing** – Within rules, verify that the fields working properly, and that the field information is populated in the database

- Trigger Processing – Within rules, verify that the trigger settings working as intended, and status information is populated in the database
- API Verification – Check for API timeouts that might be affecting the GUI

RELATED DOCUMENTATION

| [Sensors Overview | 429](#)

Use the Self Test Tool

After you set up the basic functionality in Paragon Insights, it can be challenging to diagnose problems related to configuring devices, adding devices, and applying playbooks. When an issue occurs, there are many areas that you must investigate. The self-test tool essentially acts as a fully working setup, running entirely within the Paragon Insights system.

You can refer to the following usage notes to get a better understanding of the features in the self-test tool.

- Currently this feature supports simulating devices to stream data for OpenConfig telemetry and iAgent (NETCONF).
- You can retain the self-test instance to act as a ‘running reference’.
- Do not use the self-test tool when there are undeployed changes, as the self-test tool issues its own deploy during execution.
- Do not use rules, playbooks, devices, device-groups or other elements created by the self-test tool with real network devices.

To use the self-test tool:

1. Navigate to the **Configuration > Data Ingest > Diagnostics** page, and click the **Application** tab. The Diagnostics page appears.
2. From the Self Test Settings tabbed page, configure the following self-test settings:
 - a. Select one or more sensor types from the **Sensor Type(s)** drop-down list.
 - b. (Optional) Click the **Retain Test Topology** toggle to turn it on or off.
 - If you turn the **Retain Test Topology** toggle on (default), the self-test device and device groups are retained in the **Device** and **Device Group Configuration** pages, respectively.

- If you turn the **Retain Test Topology** toggle off, the self-test resources are auto-deleted when the test is completed.

3. (Optional) Enable the **Retain Test Topology** button.

If you enable the Retain Test Topology button, the self-test device and device groups are retained in Device page and Device Group Configuration page, respectively. If you disable the button, the self-test resources are auto-deleted when the test is completes.

4. Click **Test** to run the self-test. The test results appear after a few minutes.

The color coding for the test results is as follows:

- Green status = pass
- Yellow status = error (unable to test)
- Red status = fail (test failed)
- Yellow or red status includes a message with details about the issue.

The screenshot shows the 'Debug' interface with a sidebar on the left containing 'Application', 'Reachability', 'Ingest', and 'No Data'. The main content area is titled 'Self Test Settings' and includes a 'Sensor Type(s) *' field with 'open-config' and 'iAgent' selected, and a 'Retain Test Topology' toggle set to 'Yes'. A blue 'Test' button is located at the bottom right of the settings section. Below the settings, the 'Last Test Result (Wed 09 Jun, 14:03)' is displayed with a 'hide last result' link. The 'Result Details' section is expanded, showing a table of test results for 'Ingest'.

Ingest	Test Results				
open-config	✓ Adding a device ✓ Receiving data	✓ Adding a device group	✓ Creating rule	✓ Creating playbook	✓ Deploying playbook
iAgent	✓ Adding a device ✓ Receiving data	✓ Adding a device group	✓ Creating rule	✓ Creating playbook	✓ Deploying playbook

A notification at the top right of the results table states 'Creating playbook is successful'.

The example above shows that OpenConfig and iAgent sensors are working as expected.

RELATED DOCUMENTATION

[Use the Ingest Test Tool | 567](#)

[Use the No-Data Tool | 568](#)

Use the Reachability Test

You can use the device reachability tool to diagnose device reachability issues during onboarding and any time after the onboarding process. The device reachability tool can verify connectivity to a device by using ping and SSH.

To run the reachability test:

1. Navigate to the **Configuration > Data Ingest > Diagnostics** page.
2. On the Diagnostics page, click **Reachability** to view the Reachability Test tabbed page.
3. From the Reachability Test tabbed page, configure the following settings:
 - a. Select a device group from the **Device Group** drop-down list.
 - b. Select a device from the **Device** drop-down list.
4. Click **Test** to run the device reachability test. The test results of the ping test and the SSH test appear after a few minutes.

The color coding for the test results is as follows:

- Green status = pass
- Yellow status = error (unable to test)
- Red status = fail (test failed)
- Yellow or red status includes a message with details about the issue.

Configuration / Sensor / Diagnostics

Debug

Application

Reachability

Ingest

No Data

Reachability Test

Device Group: All

Device *: vmx101

Test

hide last result

Last Test Result (Wed 09 Jun, 14:15)

Device Name	Ping	SSH
vmx101	✓PASS	✓PASS

The example above shows that the ping and the SSH tests were successful.

RELATED DOCUMENTATION

[Use the No-Data Tool](#) | 568

Use the Ingest Test Tool

The ingest connectivity tool can verify ingest methods such as OpenConfig, iAgent, and SNMP, where Paragon Insights initiates the connection with the device. The tool helps you identify a missing configuration on the network device. The tool also helps you choose appropriate playbooks and rules (that use sensors) that are compatible with the supported ingest methods.

To run the ingest connectivity test:

1. Navigate to the **Configuration > Data Ingest > Diagnostics** page.
2. On the Diagnostics page, click **Ingest** to view the Ingest Test Settings tabbed page.
3. From the Ingest Test Settings tabbed page, configure the following settings:
 - a. Select one or more sensor types from the **Sensor Type(s)** drop-down list.
 - b. Select a device group from the **Device Group** drop-down list.
 - c. Select a device from the **Device** drop-down list.
4. Click **Test** to run the ingest connectivity test. The test results appear after a few minutes.

The color coding for the test results is as follows:

- Green status = pass
- Yellow status = error (unable to test)
- Red status = fail (test failed)
- Yellow or red status includes a message with details about the issue.

The screenshot shows the 'Ingest Test Settings' page in a web application. The breadcrumb navigation is 'Configuration > Sensor > Diagnostics'. The page title is 'Debug'. On the left, there is a sidebar with 'Ingest' selected. The main content area has three fields: 'Sensor Type(s) *' with selected options 'open-config', 'iAgent', and 'snmp'; 'Device Group' with 'dg1' selected; and 'Device *' with 'vmx101' selected. A blue 'Test' button is at the bottom right. Below the settings, there is a section for 'Last Test Result (Wed 09 Jun, 14:26)'. A tooltip shows the error message: 'open-config FAIL for vmx101. DETAILS: gRPC connection failed due to login check failure'. At the bottom, a status bar for 'vmx101' shows 'open-config' with a red error icon, 'iAgent' with a green checkmark, and 'snmp' with a green checkmark. A 'hide last result' link is visible on the right.

The example above shows that iAgent and SNMP ingests are working well but gRPC connection from OpenConfig ingest failed due to login check errors.

RELATED DOCUMENTATION

[Use the Self Test Tool | 564](#)

Use the No-Data Tool

You can use the debug no-data tool to determine why a device or rule displays a “no-data” status.

These usage notes describe the features in the no-data tool.

- The tool runs through the entire sequence of checks, regardless of any issues along the way.
- The test results provide root cause information and advise where to focus your troubleshooting efforts.
- While this feature is generally intended to debug a device when it is marked as no-data, you can use it any time to verify that deployed rules are receiving data.
- This tool does not support rules using a syslog sensor, as the sensor data is event driven and not periodic.

To run the test:

1. Access the debug no-data tool by navigating to either of the following pages:

- **Monitoring > Network Health**

On this page, click any tile that shows "no-data."

Figure 45: No-Data Tile



- **Configuration > Data Ingest > Diagnostics**

The Diagnostics page appears.

2. Click **No Data** to view the No-Data Settings tabbed page.
3. From the No-Data Settings tabbed page, configure the following settings:
 - a. Select a device group from the **Device Group** drop-down list.

- b. Select a device from the **Device** drop-down list.
 - c. Select a topic from the **Topic** drop-down list.
 - d. Select one or more rules from the **Rule(s)** drop-down list.
4. Click **Test** to run the debug no-data test. The test results appear after a few minutes.
The color coding for the test results is as follows:

- Green status = pass
- Yellow status = error (unable to test)
- Red status = fail (test failed)
- Yellow or red status includes a message with details about the issue.

The screenshot shows the 'Debug' interface with the following configuration:

- Application: No-Data Settings
- Device Group: tm-dg-tenant1-juniper
- Device: vms107
- Topic: system.cpu
- Rule(s): check-system-cpu-load-average

The test results are as follows:

Component	Status
Paragon Insights Services	Fail (Red)
Device Reachability	Pass (Green)
Ingest Connectivity	Pass (Green)
Data Streaming	Pass (Green)
Field Processing	Pass (Green)
Trigger Processing	Pass (Green)
API Verification	Pass (Green)

The example above shows that all devices are working as expected.

RELATED DOCUMENTATION

[About the Network Health Page](#) | 833

Paragon Insights Tagging Overview

IN THIS SECTION

- [Tagging Profile Terminology | 570](#)
- [How Do Tagging Profiles Work? | 575](#)
- [Caveats | 576](#)

You can use the Paragon Automation graphical user interface (GUI) to create tagging profiles. You can configure a tagging profile to insert fields, values, and keys into a Paragon Insights rule. You can also set conditions that are checked against values stored in the times series database (TSDB) or the Redis database.

Paragon Insights supports the following types of tagging:

- **Static Tagging**

In static tagging, the tagging profile is applied to values stored in the time series database (TSDB). These values do not vary a lot with time. In static tagging, you can avoid using *When* statements, and you can add *Then* statements individually to a row of the TSDB. You can add tags to all rows since no conditions are present.

- **Dynamic Tagging**

In dynamic tagging, conditions used in Paragon Insights tagging are checked against values that are stored in Redis database. This database acts like a cache memory that stores dynamic data. Dynamic data is real-time data that is stored in Redis database.

For more information on tagging, see ["Types of Tagging" on page 577](#).

Tagging Profile Terminology

The following list describes the tagging profile terminologies:

Policy

A policy is the top-level element in a tagging profile. You can add multiple policies within a single tagging profile. Multiple policies that exist within a tagging profile can have their own rules and terms.

Usage Notes:

- Defining multiple policies within a single profile allows you to define terms for each rule in one profile instead of creating one profile for each rule.

Rules

A rule is any defined Paragon Insights rule. The rule element type in a tagging profile is a list element. You can apply a specific policy profile to the rule(s) (*[rule1, rule2]*) included within the tagging profile.

Usage Notes:

You can describe the topic-name/rule-name requirement for the rules element in the following ways:

- To name specific rules within a tagging profile, use the form: *topic-name/rule-name*.

For example, *protocol.bgp/check-bgp-advertised-routes*. Navigate to **Configuration>Rules** to view configured rules.

- Use an asterisk (*) with no other value or brackets to match all rules.
- Use python-based fnmatch patterns to select all rules within a specific topic. For example, *line-cards/**.

For more information, see [fnmatch – Unix filename pattern matching](#).

Terms

The term section of the tagging profile is where the match conditions are set and examined, and actions based on those matches are set and carried out. Set the conditions for a match in a *when statement*. Set the actions to be carried out upon completing a match in one or more *then statements*.

Usage Notes:

- Each term can contain a *when* statement but it is not mandatory.
- Each term must contain at least one *then* statement.
- Multiple terms can be set within a single policy.
- Terms are processed sequentially from top to bottom until a match is found. If a match is found, processing stops after completing the statements found in the *then* section. Other terms, if present, are not processed unless the *next* flag is enabled within the matched term. If the matched term has the next flag enabled, then subsequent terms are processed in order.

When Statements

When statements define the match conditions that you specify. *When* statements ultimately resolve to be true or false. You can define a *term* without a *when* statement. This equates to a default *term*

wherein the match is assumed true and the subsequent *then* statement is carried out. Conversely, multiple conditions can be checked within one *when* statement.

If one or more of the conditions set forth in a *when* statement are not met, the statement is false and the *term* has failed to match; processing moves to the next *term*, if present.

Usage Notes:

When statements perform Boolean operations on the received data to determine if it matches the criteria you set. The supported operations are:

- Numeric Operations:
 - equal-to
 - not-equal-to
 - greater-than
 - greater-than-or-equal-to
 - less-than
 - less-than-or-equal-to
- String Operations:
 - matches-with
 - does-not-match-with
- Time Operations:
 - matches-with-scheduler

NOTE: The `matches-with-scheduler` option requires that a discreet scheduler be configured in the **Administration > Ingest Settings > Scheduler** page. The name of the scheduler can then be used in the `matches-with-scheduler` *when* statement

- Go Language Expressions:
 - `eval <simple-go-expression>`

For example: `eval a + b <= c.`

Then Statements

Then statements implement the tagging instructions that you provide. This is done only after there is a complete match of the conditions set forth in a *when* statement contained in the same *term*. Each *term* defined must have at least one *then* statement. Each *then* statement must have one or more than one action(s) defined; the actions available in *then* statements are:

add-field Adds a normal field to the rule(s) listed in the rule section.

Multiple fields can be added within a *then* statement. The add-field action requires that you also define the kind of field you are adding with the *field-type* parameter:

- string
- integer
- float

NOTE: If you do not define a field type, the new field gets added with the default field-type of string.

add-key Adds a key field with string data type to the rule(s). Added key fields are indexed and can be searched for just like any other key field.

Usage Notes:

- You can set the *next* flag to true within a then statement. When this flag is set to true, the next term in the policy gets evaluated if all of the conditions of the current term match.

Example Configuration: Elements of a Tagging Profile

Paragon Insights configuration elements are displayed as pseudo-config. This configuration resembles the hierarchical method used by Junos OS.

This **Elements of a Tagging Profile** table shows how tagging profile elements are named and how they are related to each other:

```
healthbot {
  ingest-settings {
    data-enrichment {
      tagging-profile <tagging-profile-name> {
        policy <policy-name> {
          rules [ List of Rules ];
        }
      }
    }
  }
}
```


How Do Tagging Profiles Work?

You can use tagging profiles to set the conditions, define new fields and keys, and insert values into those fields. Tagging profiles are applied as part of ingest settings to allow the tags to be added to the incoming data before Paragon Insights processes the data. Since one or more rules are defined within each profile, the rules are added to a playbook and applied to a device group when the tagging profile is applied to a device.

[Table 89 on page 575](#) shows an example application identification scenario based on source-port, destination-port, and protocol of traffic seen in a NetFlow stream.

Table 89: Fields in NetFlow Stream

source-port	destination-port	protocol	derived-application
2541	Any	6 (TCP)	NetChat
Any	2541	6 (TCP)	NetChat
1755	Any	17 (UDP)	MS-streaming
Any	830	6 (TCP)	netconf-ssh
7802	Any	17 (UDP)	vns-tp

In [Table 89 on page 575](#), you use three existing fields in a NetFlow stream to guess the application traffic in the stream. You then create a new field called *derived-application* and populate it based on the values seen in the traffic.

You can apply tagging profiles at the device group level.

- When a device in a device group has a tagging profile applied to it, and the device group has another tagging profile applied to the whole group of devices, the tagging profile of the device group is merged with the existing tagging profile of the device.

For example, D-A-Net is a device that is part of a device group called Group-D1. D-A-Net has a tagging profile applied to it. There is another tagging profile applied on the device group, Group-D1, as well. In such a scenario, the tagging profile applied to the device group is merged with the tagging profile of the device, D-A-Net.

- When the tagging profile applied to the device group and the tagging profile applied to the device in the group renders the same output, the tagging profile of the device is preserved.

Example pseudo-configuration shown below

```

device r0 {
    host r0;
    tagging-profile [ profile1 ]
}
device r1 {
    host r1
}
device-group core {
    devices [ r0 r1 ];
    tagging-profile [ profile2 ]
}

```

In this example, device *r0* has tagging profile, *profile1*, assigned at the device level and tagging profile, *profile2*, assigned by its membership in the device- group (*core*).

Device *r1* has tagging profile, *profile2*, assigned by its membership in device group, *core*.

In this scenario, *profile1* and *profile2* are merged on device *r0*. However, if *profile1* and *profile2* both define the same fields but the fields contain different values, the value from *profile1* takes precedence because it is assigned directly to the device.

Device *r1* only gets tagging profile *profile2*.

Caveats

- Fields and keys added using tagging profiles cannot be used within periodic aggregation fields. This is because periodic aggregation must take place before any UDFCode function (reference, vector, UDF, ML) is applied.
- Tagging profiles can consist of only fields in add-key or add-field. Vectors cannot be added to a rule by a tagging profile.
- Vector comparison operations cannot be used within tagging profile terms. Only field Boolean operations are permitted.
- For tagging profile conditional operations within a *when* statement, the used field must be of type sensor, constant, or reference.

This is applicable only in static tagging.

- If the field used within tagging profile Boolean operation is of type reference, then this reference field must not depend on any user-defined-function or formula defined within the same rule.

RELATED DOCUMENTATION

[Types of Tagging | 577](#)

[Add a Tagging Profile | 585](#)

[Apply a Tagging Profile | 590](#)

Types of Tagging

IN THIS SECTION

- [Static Tagging | 577](#)
- [Dynamic Tagging | 580](#)

Paragon Insights supports static tagging and dynamic tagging.

Static Tagging

In static tagging, the tagging profile is applied to values stored in the time series data base (TSDB). These values do not vary a lot with time. In static tagging, you can avoid using *When* statements, and you can add *Then* statements to a tagging profile.

Sample Static Tagging Configuration

```
healthbot {
  ingest-settings {
    data-enrichment {
      tagging-profile profile {
        policy policy1 {
          rules *;
          term term1 {
            then {
              add-key "tenant-id" {
                value tenant1;
              }
            }
          }
        }
      }
    }
  }
}
```

```

    }
  }
}
}
}

```

In this sample static tagging configuration, the lack of a *when* statement means that any device that this tagging profile is applied to will have the field *tenant-id* assigned with the value *tenant1*. The fields and values defined in this profile are assigned to all rules that are applied to a device or device-group because of the *** in the rules parameter.

You can also create a static tagging profile from the Paragon Automation graphical user interface (GUI). Navigate to **Configuration > Sensor > Settings > Tagging Profile** page to create a tagging profile.

Application Identification

[Table 90 on page 578](#) shows an example application identification scenario based on source-port, destination-port, and protocol of traffic seen in a NetFlow stream.

Table 90: Fields in NetFlow Stream

source-port	destination-port	protocol	derived-application
2541	Any	6 (TCP)	NetChat
Any	2541	6 (TCP)	NetChat
1755	Any	17 (UDP)	MS-streaming
Any	830	6 (TCP)	netconf-ssh
7802	Any	17 (UDP)	vns-tp

To create the derived-application field as given in [Table 90 on page 578](#) from the received data (data under source-port, destination port, and protocol), you must use a tagging profile definition that looks like this:

```

healthbot {
  ingest-settings {

```

```
data-enrichment {
  tagging-profile profile1 {
    policy policy1 {
      rules *;
      term term1 {
        when {
          matches-with "$source-port" "$netchat-source-port";
          matches-with "$protocol" "6 (TCP)";
        }
        then {
          add-key "application" {
            value netchat;
          }
        }
      }
      term term2 {
        when {
          matches-with "$protocol" "6 (TCP)";
          matches-with "$destination-port" "$netchat-dest-port";
        }
        then {
          add-key "application" {
            value netchat;
          }
        }
      }
      term term3 {
        when {
          matches-with "$source-port" "$ms-streaming-source-port";
          matches-with "$protocol" "17 (UDP)";
        }
        then {
          add-key "application" {
            value ms-streaming;
          }
        }
      }
      term term4 {
        when {
          matches-with "$source-port" "$netconf-ssh-source-port";
          matches-with "$protocol" "6 (TCP)";
        }
        then {
```


- Device Group

Core:r1::/components/

- Network Group

network::net_check::topic/rule

- Values are stored in JSON string format `<json dump as string>` in Redis. However, values are provided in string, integer, and float formats.

Example value formats:

- Core:r1::/components/= value1
- Core:r1::/components/='{“key1”: value1, “key2”: value2}'
- Core:r1::/components/='{“key1”: {“key2”: value1, “key3”: value2}'
- Core:r1::/components/='{“key1”: {“key2”: “[list of values]”, “key3”: value1}'
- Sample tagging-profile configurations using when statement.

```
"when" : {
    "matches-with" : [
        {
            "left-operand" : "$field1",
            "right-operand" : "/interfaces/.key1",
            "in-memory": true
        }
    ]
}
```

- Use a . operator between interfaces.

In the following example, key3 interface is nested within key2 interface in the right operand.

```
"when" : {
    "matches-with" : [
        {
            "left-operand" : "$field1",
            "right-operand" : "/interfaces/.key2.key3",
        }
    ]
}
```

```

    ]
  }

```

- Sample tagging-profile configurations using then statement.

```

"then" : {
  "add-field" : [
    {
      "name" : "field1",
      "value" : "redis-key",
      "type" : "integer",
      "in-memory": true
    }
  ]
}

```

- Use a . operator between interfaces.

In the following example, key3 interface is nested within key2 interface in the right operand.

```

"then" : {
  "add-field" : [
    {
      "name" : "field2",
      "value" : "redis-key1.redis-key2.redis-key3",
      "type" : "integer",
      "in-memory": true
    }
  ]
}

```

- Using exist operator in configurations.
- Using exist as key.

Redis Data Structure

```

“Core:r1::/interfaces/” = ‘{“ge-1/0/2”: {“key1”: value1, “key2”: value2}’

```

tagging-profile Using when Statement

```

“when”: {
  “exists”: {
    “field”: “$interface-name”,
    “path”: “/interfaces/”,
    “in-memory”: true
  }
}
“then”: { do-something.. }

```

- Using exist as value in list.

Redis Data Structure

```

“Core:r1::/interfaces/” = ‘{“key1”: {“key2”: [‘ge-1/0/2’, ‘ge-1/0/3’], “key3”: value1}}’

```

tagging-profile Using when Statement

```

“when”: {
  “exists”: {
    “field”: “$interface-name”,
    “path”: “/interfaces/.key1.key2” ,
    “in-memory”: true
  }
},
“then”: {
  “add-field”: [
    “name”: “field1”,
    “value”: “/interfaces/.key1.key3”,
    “in-memory”: true
  ]
}

```

- Using \$ in **then** statements.

When you use \$<field-name> within a Redis key, \$<field-name> is replaced with a value from the already processed database value.

As an example, consider that ge-1/0/2 is present within Redis key.

Redis Data Structure

```
“Core:r1::/interfaces/” = ‘{“ge-1/0/2”: {“key1”: value1, “key2”: value2},
  “ge-1/0/3”: {“key1”: value1, “key2”: value2}’
```

Example tagging -profile

```
“when”: {
  “exists”: [
    {
      “field”: “$interface-name”,
      “path”: “/interfaces/”,
      “in-memory”: true
    }
  ],
  “greater-than”: [
    {
      “left-operand”: “30”,
      “right-operand”: “/interfaces/.$interface-name.key1” ,
      “in-memory”: true
    }
  ]
},
“then”: {
  “add-field”: [
    “name”: “interface-meta-data”,
    “value”: “/interfaces/.$interface-name.key2”,
    “in-memory”: true
  ]
}
```

In this scenario, the tagging-profile checks if \$interface-name is present in the Redis database, and if key1 value for the given interface name is greater than 30. If the statement is true, the tagging-profile fetches key2 value from name field. In this example tagging profile, the name value is interface-meta-data.

- To enable dynamic tagging, set in-memory value to true.

By default in-memory value is set to false.

```
“when”: {
  “exists”: {
```

```

“field”: “$interface-name”,
  “path”: “/interfaces/.key1.key2” ,
  “in-memory”: true
}
}
“then”: {
“add-field”: [
“name”: “interface-meta-data”,
  “value”: “/interfaces/.$interface-name.key2”,
  “in-memory”: true
]
}
}

```

RELATED DOCUMENTATION

[Paragon Insights Tagging Overview | 570](#)

[Add a Tagging Profile | 585](#)

[Apply a Tagging Profile | 590](#)

Add a Tagging Profile

You can use the Paragon Automation graphical user interface (GUI) to add static tagging and dynamic tagging profiles.

For more information on tagging, see "[Paragon Insights Tagging Overview](#)" on page 570.

Adding a Static Tagging Profile

To add a static tagging profile:

1. Navigate to **Configuration > Data Ingest > Settings**.

The **Ingest Settings** page is displayed.

2. Click the **Tagging Profile** tab and then click the plus (+) icon to add a tagging profile.

The **Create Tagging Profile** page is displayed.

3. Enter the following information in the **Create Tagging Profile** page:

- a. Enter a name for the tagging profile in the **Profile Name** text box.

The maximum length is 64 characters.

Regex pattern: “[a-zA-Z][a-zA-Z0-9_-]*”

- b. Click the plus (+) icon under Policies to define a policy for this tagging profile.

You can define one or more policies.

The **Policies** section is displayed.

- i. Enter a name for the new policy in the **Policy Name** text box.

The maximum length is 64 characters.

Regex pattern: “[a-zA-Z][a-zA-Z0-9_-]*”

- ii. Enter a rule that you want to apply to this tagging profile. The rule can contain an *fnmatch* expression.

You can apply one or more than one rule to a profile. A rule is any defined Paragon Insights rule.

- c. Click the plus (+) icon under **Terms** to define a list of conditions.

- i. Enter a name for the match condition in the **Term Name** text box.

- ii. Configure When and Then statements.

You can define tagging instructions in a Then statement. After the conditions that you set in a When statement are met, the Then statement is implemented. When statements are mandatory in static tagging.

To configure a **Then** statement:

1. Click the plus (+) icon to add a key to the rules listed.

The **Key Name** and **Value** text boxes are displayed.

- Enter a name for the key in the **Key Name** text box.

The maximum length is 64 characters.

Regex pattern: “[a-zA-Z][a-zA-Z0-9_-]*”

This name is added as the key field for all rules configured within the tagging profile rules section.

- Enter a value that you want to associate to the key, in the **Value** text box.

2. Click the plus (+) icon to add a field to the rules listed.

The **Field Name** and **Value** text boxes, and the **Type** drop-down list are displayed.

- Enter the name in the **Field Name** text box.

- Enter a value in the **Value** text box.

- Select the field type from the **Type** drop-down list.

String type is selected by default.

You can also select unsigned integer as a name field data type. An unsigned integer is a data type that can contain values from 0 through 4,294,967,295.

- Click **OK**.

3. Set the **Evaluate next term** flag to **True** to evaluate conditions in the next term. Evaluate next term only if the first condition is satisfied.

By default, the **Evaluate next term** flag is set to **False**.

4. Click **Save** to only save the configuration.

Click **Save & Deploy** to save and deploy the configuration immediately.

Adding a Dynamic Tagging Profile

To configure a dynamic tagging profile with Redis:

1. Navigate to **Configuration > Data Ingest > Settings**.

The **Ingest Settings** page appears.

2. Click the **Tagging Profile** tab and then click the plus (+) icon to add a tagging profile.

The **Create Tagging Profile** page is displayed.

3. Enter the following information in the **Create Tagging Profile** page.

- a. Enter a name for the tagging profile in the **Profile Name** text box.

The maximum length is 64 characters.

Regex pattern: “[a-zA-Z][a-zA-Z0-9_-]*”

- b. Click the plus (+) icon under **Policies** to define a policy for this tagging profile.

You can define one or more policies.

The **Policies** section is displayed.

- i. Enter a name for the new policy in the **Policy Name** text box.

The maximum length is 64 characters.

Regex pattern: “[a-zA-Z][a-zA-Z0-9_-]*”

- ii. Enter a rule that you want to apply to this tagging profile. The rule can contain an *fnmatch* expression.

You can apply one or more rules to a profile. A rule is any defined Paragon Insights rule.

- c. Click the plus (+) icon under **Terms** to define a list of conditions.
 - i. Enter a name for the match condition in the **Term Name** text box.
 - ii. Configure When and Then statements:

You set conditions for a match in a when statement. To configure **When** statement,

1. Click the **Edit (pencil)** icon.

The **When Condition** page is displayed.

2. Click **+ Add another when** to view the **Operator** drop-down list.
3. Select a boolean operation that you want to apply to incoming data from the **Operator** drop-down list.

The **Left Operand** and **Right Operand** text boxes are displayed.

NOTE: The **When Condition** drop-down list is automatically renamed to the operator condition that you selected.

- Enter the value of the left operand of assignment that you selected, in the **Left Operand** text box.

You can use \$ as prefix to populate database values. For example, \$memory. However, using \$ as prefix is not mandatory.

- Enter the value of the right operand of assignment that you selected, in the **Right Operand** text box.

This value is populated from the Redis database.

- Set the **Evaluate in Memory** flag to **True** to populate data from the Redis database.

By default, the **Evaluate in Memory** flag is set to **False**. When the flag is set to false, data is populated from the TSDB.

- Click **OK**.

4. Set the **Evaluate next term** flag to **True** to evaluate conditions in the next term. After the first condition is met, the conditions in the next term are evaluated.

By default, the **Evaluate next term** flag is set to **False**.

You can define tagging instructions in a Then statement. After the conditions that you set in a When statement are met, the Then statement is implemented. When statements are mandatory.

To configure a **Then** statement:

1. Click the plus (+) icon to add a key to the rules listed.

The **Key Name** and **Value** text boxes are displayed.

- Enter a name for the key in the **Key Name** text box.

The maximum length is 64 characters.

Regex pattern: “[a-zA-Z][a-zA-Z0-9_-]*”

This name will be added as key field for all rules configured within the tagging profile rules section.

- Enter a value that you want to associate with the key, in the **Value** text box.

2. Click the plus (+) icon to add a text box to the rules listed.

The **Field Name** and **Value** text boxes, and the **Type** drop-down list are displayed.

- Enter the name in the **Field Name** text box.
- Enter a value in the **Value** text box.
- Select the field type from the **Type** drop-down list.

String type is selected by default.

You can also select unsigned integer as a name field data type. An unsigned integer is a data type that can contain values from 0 through 4,294,967,295.

- Set the **Evaluate in Memory** flag to **True** to populate data from the Redis database.

By default, the **Evaluate in Memory** flag is set to **False**.

- Click **OK**.

3. Set the **Evaluate next term** flag to **True** to evaluate conditions in the next term. The next term is evaluated only if the first condition is satisfied.

By default, the **Evaluate next term** flag is set to **False**.

4. Click **Save** to only save the configuration.

Click **Save & Deploy** to save and deploy the configuration immediately.

RELATED DOCUMENTATION

[Paragon Insights Tagging Overview | 570](#)

[Types of Tagging | 577](#)

[Apply a Tagging Profile | 590](#)

Apply a Tagging Profile

You can configure a tagging profile to insert fields, values, and keys into a Paragon Insights rule. You can also set conditions that are checked against values stored in the times series database (TSDB) or Redis database. For more information on configuring a tagging profile, see ["Add a Tagging Profile" on page 585](#).

After you have created a tagging profile from the Paragon Automation graphical user interface (GUI), you can apply a tagging profile to:

- a new device group
- an existing device group

Follow these steps to apply a tagging profile.

To apply a tagging profile to a new device group:

1. Navigate to **Configuration > Device Groups**.

The **Device Group Configuration** page is displayed.

2. Click (+) icon to add a new device group.

The **Add Device Group** page is displayed.

3. Enter the following information in the Add Device Group page.

To create a new device group:

- a. Enter a name for the device group in the **Name** text box.
- b. Enter a description for the device group in the **Description** text box.
- c. Select the devices from the **Devices** drop-down list that you want to add to the device group.
- d. Click the **Tagging Profiles** section and select the tagging profile you want to apply to the device, from the **Profiles** drop-down list
- e. Click **Save** to only save the configuration.
Click **Save & Deploy** to save and deploy the configuration immediately.

For more information on adding a device group, see ["Add a Device Group" on page 177](#).

To apply a tagging profile to an existing device group:

1. Navigate to **Configuration > Device Groups**.

The **Device Group Configuration** page is displayed.

2. Select the check box next to the name of the device group and click the **Edit device group** icon.

The **Edit Device Group** page is displayed.

3. Click the **Advanced > Tagging Profiles** section to view the **Profiles** drop-down list.

4. Select the tagging profile you want to apply to the device group, from the **Profiles** drop-down list.

5. Click **Save** to only save the configuration.

Click **Save & Deploy** to save and immediately deploy the configuration.

For more information on adding a device group, see "[Add a Device Group](#)" on page 177.

RELATED DOCUMENTATION

[Paragon Insights Tagging Overview | 570](#)

[Types of Tagging | 577](#)

[Add a Tagging Profile | 585](#)

Delete a Tagging Profile

To delete a Tagging Profile:

1. Click **Configuration > Data Ingest > Settings** from the left-nav bar.

The Ingest Settings page is displayed.


2. Click the **Tagging Profile** tab to view the Tagging Profile page.

3. Select the tagging profile that you want to delete, and click the **delete (trash can)** icon.

The **CONFIRM DELETE** pop-up appears.

4. Do one of the following:

Figure 46: Confirm Delete Pop-up



- Click **Yes** to delete the profile from the database. However, the changes are not applied to the ingest service.

NOTE:

- We recommended that you do not delete a tagging profile that is currently in use.
 - After you delete a tagging profile from the database, you cannot associate that tagging profile with another device or device group even if you have not deployed changes.
 - You can also deploy changes to the ingest service or roll back the changes that you have already deleted, from the **Health Configuration Deployment Status** page. For more information, see "[Commit or Roll Back Configuration Changes in Paragon Insights](#)" on [page 185](#).
- Select the **Deploy changes** check box and then click **Yes** to delete the profile from the database, and to apply the changes to the ingest service.
 - (Optional) Click **No** to cancel this operation.

The Tagging Profile is deleted.

Understand User-Defined Actions and Functions

When creating rules, Paragon Automation Platform includes the ability to run user-defined actions (UDAs) as part of a trigger. A Rule executes UDAs as Python scripts. For example, you might configure a rule with a trigger that reacts to some critical interface going down and responds to the event by calling a function to send an SMS alert. You can write the logic to send the SMS in a UDA python script.

In Paragon Automation, you can schedule UDAs and notifications. This is useful when you deploy multiple parallel instances of Paragon Automation Platform in different locations. You can schedule UDAs to run alternatively from UDA schedulers located in different regions. In the event of a node failure, the UDA scheduler running in a parallel instance continues to execute your UDA and notifications. To enable UDA scheduler, see ["Enable UDA Scheduler in Trigger Action" on page 599](#).

You can also includes the ability to run user-defined functions (UDFs). Also created as Python scripts, UDFs provide the ability to process incoming telemetry data from a device and store the processed value in the time-series database for that device/device-group. For example, the device may be sending FPC temperature in Celsius but you want to process it to be stored as Fahrenheit values in the database.

The processing of UDF fields is handled by microservices called UDF farm unless, you declare global variables in Python scripts. This approach allows for Paragon Insights to process multiple data points from multiple devices and fields at the same time (parallel processing). The result is a 4 to 5 times increase in processing performance for UDA/UDF.

Global Variables in Python Scripts

TAND executes Python scripts that use global variables. Global variables retain a value across multiple UDFs.

The following is an example function to calculate cumulative sum and store the value in global variable *sum*.

```
sum = 0
def cumulative_sum (a, b):
    global sum
    sum = sum + a + b
    return sum
```

When you use global variables in Python scripts, the UDFs are processed by TAND instead of UDF farms.

As an alternative to global variables, you can use the Python construct `**kwargs` to capture values that must be retained across different functions. When Paragon Insights calls a function (defined in a UDF), it sends topic name, rule name, device group, point time, and device ID that are captured using the

construct `**kwargs`. In case of UDAs, Paragon Automation sends topic name and rule name while executing the Python script.

Along with infrastructure values, Paragon Insights also sends a parameter called `hb_store` in `**kwargs` that fetches the last computed value for a variable.

To illustrate how `hb_store` works in the cumulative addition example:

```
def sum(a, b, **kwargs):
    if 'sum' not in kwargs[hb_store]:
        kwargs[hb_store]['sum'] = 0 #if 'sum' is not present in kwargs, declare the
initial 'sum' value as 0.

        kwargs[hb_store]['sum'] = kwargs[hb_store]['sum'] + a + b #Store cumulative
addition value in 'sum'

    return kwargs[hb_store]['sum']
```

Each time a function with the above code is called, it performs addition of last stored value in 'sum' with the value of *a* and value of *b*. The new value of addition operation is displayed and stored in 'sum'.

RELATED DOCUMENTATION

| [Add a Predefined Rule](#) | 340

Modify User-Defined Action, Function, and Workflow Engines

The following section describes how you can modify the UDA/UDF/workflow engine in Paragon Automation command line interface (CLI).

You can modify the UDA, UDF, or workflow engine using the Paragon Automation CLI, as shown below.

NOTE: You must run the following bash commands from the primary node of Paragon Automation.

```

user@paragon-master:/var/local/healthbot# ./healthbot modify-uda-engine --help
usage: healthbot modify-uda-engine [-h] (-s SCRIPT | --rollback) [--simulate]

optional arguments:
  -h, --help            show this help message and exit
  -s SCRIPT, --script SCRIPT
                        Run script in UDA engine
  --rollback, -r        Rollback UDA engine to original state
  --simulate            Run script in simulated UDA engine and show output

user@paragon-master:/var/local/healthbot# ./healthbot modify-udf-engine --help
usage: healthbot modify-udf-engine [-h] (-s SCRIPT | --rollback) [--simulate]
                                     [--service SERVICE]

optional arguments:
  -h, --help            show this help message and exit
  -s SCRIPT, --script SCRIPT
                        Run script in UDF engine
  --rollback, -r        Rollback UDF engine to original state
  --simulate            Run script in simulated UDF engine and show output
  --service SERVICE     Modify specific service UDF

root@paragon-master:/var/local/healthbot# ./healthbot modify-workflow-engine --help

usage: healthbot.py modify-workflow-engine [-h] (-s SCRIPT | --rollback)
                                           [--simulate]

optional arguments:
  -h, --help            show this help message and exit
  -s SCRIPT, --script SCRIPT
                        Run script in WORKFLOW engine
  --rollback, -r        Rollback WORKFLOW engine to original state
  --simulate            Run script in simulated WORKFLOW engine and show output

```

The commands have three main options:

- Simulate—test a script (and view its output) in the simulated UDA/UDF/workflow engine environment without affecting the running Paragon Automation system
- Modify—modify the actual UDA/UDF/workflow engine using a script
- Rollback—revert to the original version of the UDA/UDF/workflow engine

Usage Notes

- The bash script will run in a container running Ubuntu OS Release 16.04 or 18.04; write the script accordingly.
- The script must be non-interactive; any questions must be pre-answered. For example, use the '-y' option when installing a package using apt-get.
- If you prefer to copy the source packages of the dependency modules onto the Paragon Insights server so the engine can manually install them instead of downloading them from the Internet, place the required source packages in the `/var/local/healthbot/input` directory. Then within your bash script, point to the `/input` directory. For example, to use a file placed in `/var/local/healthbot/input/myfile.txt`, set the bash script to access it at `/input/myfile.txt`.
- Modifying the UDA/UDF/workflow engine more than once is *not* an incremental procedure; use a new bash script that includes both the original and new instructions, and re-run the modify procedure using the new script.
- Modifications to UDA/UDF/workflow engines are applicable in current installation.

Once you upgrade the version, you must run the script to modify UDA/UDF/workflow engines.

NOTE: The following examples use the UDA engine; these procedures apply equally to the UDF and workflow engines.

NOTE: The following procedure assumes that you installed Paragon Automation server.

Simulate

Use the simulate feature to test your bash script in the simulated environment, without affecting the running Paragon Insights system.

To simulate modifying the UDA engine:

1. Enter the command `./healthbot modify-uda-engine -s /<path>/<script-file> --simulate`.

2. The script runs and the output shows on screen, just as if you entered the script commands yourself.

```

user@paragon-master:/var/local/healthbot# ./healthbot modify-uda-engine -s /var/tmp/test-
script.sh --simulate
Running /var/tmp/test-script.sh in simulated alerta engine..
Get:1 http://security.ubuntu.com/ubuntu xenial-security InRelease [109 kB]
...
Fetched 4296 kB in 15s (278 kB/s)
Reading package lists...

Building dependency tree...

Reading state information...
...

```

Modify

When you are satisfied with the simulation results, go ahead with the actual modification procedure.

To modify the UDA engine:

1. Load the desired bash script onto the Paragon Insights server.
2. If your Paragon Insights server is fully up and running, issue the command `./healthbot stop -s alerta` to stop the running services.
3. Run the command `./healthbot modify-uda-engine -s /<path>/<script-file>`.

```

user@paragon-master:/var/local/healthbot# ./healthbot modify-uda-engine -s /var/tmp/test-
script.sh
Running /var/tmp/test-script.sh in simulated alerta engine..
Success! See /tmp/.alerta_modification.log for logs
Please restart alerta by issuing './healthbot start --device-group healthbot -s alerta'

```

4. (Optional) As noted in the output, you can check the log file to further verify the script was loaded successfully.
5. Restart the alerta service using the command `./healthbot start -s alerta`.
6. Once complete, verify that the alerta service is up and running using the command `./healthbot status`.

7. To verify that the UDA engine has been updated, use the command `./healthbot version -s alerta` and check that the `healthbot_alerta` container is using the `<version>custom` tag.

```
user@paragon-master:/var/local/healthbot# ./healthbot version -s alerta
{'alerta': 'healthbot_alerta:2.1.0-custom'}
```

The UDA engine is now running with the installed dependencies as per the bash script.

Rollback

If you have a need or desire to remove the changes to the engine, you can revert the engine to its original state.

To rollback the UDA engine:

1. Enter the command `./healthbot modify-uda-engine --rollback`.

```
user@paragon-master:/var/local/healthbot# ./healthbot modify-uda-engine --rollback
Rolling back alerta engine to original state..
Successfully rolled back alerta engine
Please restart alerta by issuing './healthbot start --device-group healthbot -s alerta'
```

Note that it is not necessary to restart the `alerta` service at this point.

2. Once complete, verify that the `alerta` service is up and running using the command `./healthbot status`.
3. To verify that the UDA engine has reverted back, use the command `./healthbot version -s alerta` and check that the `healthbot_alerta` container is using the `<version>` tag.

```
user@paragon-master:/var/local/healthbot# ./healthbot version -s alerta
{'alerta': 'healthbot_alerta:2.1.0'}
```

The UDA engine is now running in its original state, with no additional installed dependencies.

RELATED DOCUMENTATION

| [Understand User-Defined Actions and Functions](#) | 593

Enable UDA Scheduler in Trigger Action

You can schedule user-defined actions (UDAs) and notifications to be executed within a set time interval. To schedule UDAs, you must first create a discrete scheduler and then link the scheduler in the Trigger Action page.

NOTE: You can link only one trigger action scheduler to a Paragon Automation Platform instance.

To know more about creating a scheduler, see ["Configure Scheduler Settings" on page 630](#).

To enable a scheduler:

1. Go to **Configuration > Insights Settings**.
2. Click the **Trigger Action** tab.
3. Select a scheduler profile that you want to associate with Trigger Action.
4. Do one of the following:
 - Click **Save** to save the scheduler profile.
Paragon Automation does not apply the profile to the device or the network groups. This option enables you to commit or rollback the configuration changes in the platform.
 - Click **Save and Deploy** to deploy the configuration in your Paragon Automation instance.
Paragon Automation executes the UDA and notifications based on the time period and the time interval configured in the scheduler. To cancel UDA scheduling, you can remove the scheduler profile and repeat the save and deploy option.

The scheduler set in Trigger Action applies to all device groups and network groups. You can disable UDA scheduling in the network group configuration. For more information, see:

- ["Add a Network Group" on page 266](#)

RELATED DOCUMENTATION

[Understand User-Defined Actions and Functions | 593](#)

Understand kube-state-metrics Service

IN THIS SECTION

- [List of Metrics that are Exposed | 600](#)
- [Enable kube-state-metrics | 601](#)
- [Sample Rules and Playbooks | 601](#)

NOTE: kube-state-metrics is supported as a beta-only feature.

kube-state-metrics is a third-party metrics monitoring service that generates metrics based on the current state of Kubernetes clusters. You can use kube-state-metrics to monitor the health of Kubernetes cluster and services. kube-state-metrics service is supported as a beta-only feature. kube-state-metrics runs as a cluster service, and is installed automatically when you install Paragon Automation. Once this service is installed, you can enable this service to generate, monitor, and expose metrics of various objects within a Kubernetes cluster.

kube-state-metrics service provides metrics on pods, DaemonSets, deployments, persistent volume, endpoints, ingress, job, lease, and configmap objects that are part of a Kubernetes cluster.

List of Metrics that are Exposed

The following is the list of metrics that are exposed:

- Pods running in a namespace
- Pods that are available
- Information on successful/failed deployments
- State of persistent volumes
- Information on currently running, successful, and failed jobs
- Pods that are in error state
- Health of deployment and DaemonSets
- Status and condition of Kubernetes nodes

Enable kube-state-metrics

You can enable kube-state-metrics from the CLI as well as from the Paragon Automation UI.

The following is an overview of steps to enable kube-state-metrics from the UI:

1. Create an unmanaged device.

The unmanaged device represents the cluster. The device hostname must be the kube-state-metrics service IP. For example, kube-state-metrics.healthbot.svc.cluster.local.

2. Add the device to a device group.

3. Create rules.

4. Apply playbooks.

Sample Rules and Playbooks

check-daemonset-status.rule

```
healthbot {
  topic kube-metrics {
    rule check-daemonset-status {
      keys [ daemonset namespace ];
      synopsis "";
      description "Checks daemon set unavailable status";
      sensor daemonset-status {
        description "Checks daemon set unavailable status";
        server-monitoring {
          sensor-name /kube/daemonset;
          frequency 60s;
        }
      }
    }
    field daemonset {
      sensor daemonset-status {
        path daemonset;
      }
      type string;
      description "Checks status for demonset key";
    }
    field daemonset_status {
      sensor daemonset-status {
        path /kube/daemonset/status/number/unavailable;
      }
    }
  }
}
```



```

rule check-deployment-status-condition {
  keys [ condition deployment namespace ];
  synopsis "";
  description "Checks kube metrics deployment status condition";
  sensor deployment-status {
    description "Checks kube metrics deployment status condition";
    server-monitoring {
      sensor-name /kube/deployment;
      frequency 60s;
    }
  }
  field condition {
    sensor deployment-status {
      path condition;
    }
    type string;
    description "Deployment condition";
  }
  field deployment {
    sensor deployment-status {
      path deployment;
    }
    type string;
    description "Deployment pod name";
  }
  field namespace {
    sensor deployment-status {
      path namespace;
    }
    type string;
    description "Deployment namespace";
  }
  field status {
    sensor deployment-status {
      path status;
    }
    type string;
    description "Checks for true or false condition";
  }
  trigger deployment-status {
    frequency 1offset;
    term available {
      when {

```



```

        description "Deployment pod name";
    }
    field deployment_status {
        sensor deployment-status {
            path /kubernetes/deployment/status/replicas;
        }
        type float;
        description "Field to check 0 or other values";
    }
    field namespace {
        sensor deployment-status {
            path namespace;
        }
        type string;
        description "Namespace key";
    }
    trigger deployment-status {
        frequency 1offset;
        term available {
            when {
                not-equal-to "$deployment_status" 0;
            }
            then {
                status {
                    color green;
                    message "Deployment status for replicaset is $deployment_status for
$namespace $deployment";
                }
            }
        }
        term notavailable {
            then {
                status {
                    color red;
                    message "Deployment status for replicaset is $deployment_status for
$namespace $deployment";
                }
            }
        }
    }
}

```

```

    }
}

```

check-node-status.rule

```

healthbot {
  topic kube-metrics {
    rule check-node-status {
      keys [ condition node ];
      synopsis "";
      description "Checks node status";
      sensor node-status {
        description "Checks node status";
        server-monitoring {
          sensor-name /kube/node;
          frequency 60s;
        }
      }
      field condition {
        sensor node-status {
          path condition;
        }
        type string;
        description "Node condition";
      }
      field node {
        sensor node-status {
          path node;
        }
        type string;
        description "Node name";
      }
      field node_status {
        constant {
          value "{{value}}";
        }
        type integer;
        description "Field to check condition";
      }
      field status {
        sensor node-status {
          path status;

```

```

    }
    type string;
    description "Status of the node";
  }
  trigger node-status {
    frequency 1offset;
    term available {
      when {
        matches-with "$status" false {
          ignore-case;
        }
      }
      then {
        status {
          color green;
          message "$condition for $node is $status";
        }
      }
    }
    term notavailable {
      then {
        status {
          color red;
          message "$condition for $node is $status.";
        }
      }
    }
  }
  variable value {
    value 0;
    description "Variable to match true(0) condition.";
    type int;
  }
}
}
}

```

check-pod-container-restarts.rule

```

healthbot {
  topic kube-metrics {
    rule check-pod-container-restarts {

```

```

keys [ container namespace pod uid ];
synopsis "";
description "Checks pod container status restarts";
sensor pod-init-container-status {
    description "Checks pod container status restarts";
    server-monitoring {
        sensor-name /kube/pod/container;
        frequency 60s;
    }
}
field container {
    sensor pod-init-container-status {
        path container;
    }
    type string;
    description "container name";
}
field namespace {
    sensor pod-init-container-status {
        path namespace;
    }
    type string;
    description "namespace of the pod";
}
field pod {
    sensor pod-init-container-status {
        path pod;
    }
    type string;
    description "pod name";
}
field restart-value {
    sensor pod-init-container-status {
        path /kube/pod/container/status/restarts/total;
    }
    type integer;
    description "restart status value";
}
field uid {
    sensor pod-init-container-status {
        path uid;
    }
    type string;
}

```

```

        description "Id ";
    }
    trigger restart-total-status {
        frequency 1offset;
        term less-than-ten {
            when {
                less-than-or-equal-to "$restart-value" 10;
            }
            then {
                status {
                    color green;
                    message "$namespace pod $pod container $container restart total is
$restart-value ";
                }
            }
        }
        term between-ten-and-twenty {
            when {
                range "$restart-value" {
                    min 10;
                    max 20;
                }
            }
            then {
                status {
                    color yellow;
                    message "$namespace pod $pod container $container restart total is
$restart-value ";
                }
            }
        }
        term more-than-twenty {
            then {
                status {
                    color red;
                    message "$namespace pod $pod container $container restart total is
$restart-value";
                }
            }
        }
    }
}

```

```

    }
}

```

check-pod-init-container-status.rule

```

healthbot {
  topic kube-metrics {
    rule check-pod-init-container-status {
      keys [ container namespace pod uid ];
      synopsis "";
      description "Checks pod init container status waiting";
      sensor pod-init-container-status {
        description "Checks pod init container status waiting";
        server-monitoring {
          sensor-name /kube/pod/init/container;
          frequency 60s;
        }
      }
    }
    field container {
      sensor pod-init-container-status {
        path container;
      }
      type string;
      description "container name";
    }
    field namespace {
      sensor pod-init-container-status {
        path namespace;
      }
      type string;
      description "namespace of the pod";
    }
    field pod {
      sensor pod-init-container-status {
        path pod;
      }
      type string;
      description "pod name";
    }
    field status {
      sensor pod-init-container-status {
        path /kube/pod/init/container/status/waiting;
      }
    }
  }
}

```


kube-metrics.playbook

```
healthbot {
  playbook kube-metrics {
    rules [ kube-metrics/check-daemonset-status kube-metrics/check-deployment-status-
replicas kube-metrics/check-node-status kube-metrics/check-pod-container-restarts kube-metrics/
check-pod-init-container-status kube-metrics/check-deployment-status-condition ];
    description "Rules to check for kube-metrics ";
    synopsis "Rules to check for kube-metrics ";
  }
}
```

Insights Settings

IN THIS CHAPTER

- [About the Insights Settings Page | 613](#)
- [Add Alert Blackouts | 617](#)
- [About Alert Notifications | 619](#)
- [Use Exim4 for E-mails | 619](#)
- [Configure the Exim4 Agent to Send E-mail | 620](#)
- [Configure a Notification Profile | 621](#)
- [Enable Alert Notifications for Device Groups and Network Groups | 627](#)
- [Configure Report Settings | 628](#)
- [Configure Scheduler Settings | 630](#)
- [Configure a Retention Policy | 632](#)
- [Configure Destination Settings | 633](#)
- [Time Series Database \(TSDB\) Overview | 635](#)
- [Manage Time Series Database Settings | 638](#)
- [Backup and Restore the TSDB | 640](#)
- [Time Series Database Replication Scenarios | 642](#)

About the Insights Settings Page

IN THIS SECTION

- [Tasks You Can Perform | 614](#)
- [Field Descriptions | 614](#)

To access this page from the Paragon Automation graphical user interface (GUI), click **Configuration > Insights Settings**.

Each time you navigate to **Configuration > Insights Settings** page, you first see the **Alert Blackout Settings** tab. To view the other tabbed pages, click the tabs on the left. For example, click **Retention Policy** to view the **Retention Policy Settings** tab.

Tasks You Can Perform

You can perform the following tasks from this page:

- Add blackout periods for an alarm. See ["Add Alert Blackouts" on page 617](#).
- Configure a notification profile. See ["Configure a Notification Profile" on page 621](#).
- Configure a retention policy for the time series database (TSDB). See ["Configure a Retention Policy" on page 632](#).
- Configure destination settings. See ["Configure Destination Settings" on page 633](#).
- Create scheduler profiles. See ["Configure Scheduler Settings" on page 630](#).
- Generate a report on demand. See ["Configure Report Settings" on page 628](#).
- Configure TSDB settings. See ["Manage Time Series Database Settings" on page 638](#).

Field Descriptions

[Table 91 on page 614](#) describes the fields on the Ingest Settings page.

Table 91: Fields on the Insights Settings Page

Field	Description
Fields on the Alert Blackout Settings page	
Device Group	View the device group to which the blackout profile is applied.
Topic	View the topic(s) to which the blackout profile is applied.

Table 91: Fields on the Insights Settings Page (Continued)

Field	Description
Resource	View the resource(s) to which the blackout profile is applied.
Event	View the event(s) to which the blackout profile is applied.
Start	View the start time of the blackout profile.
End	View the end time of the blackout profile.

Fields on the Notification Settings page**Table 91: Fields on the Insights Settings Page (Continued)**

Name	View the name of the notification profile.
Type	View the notification type attribute for a notification profile.

Table 91: Fields on the Insights Settings Page (Continued)**Fields on the Retention Policy Settings page**

Name	View the name of the retention policy.
Duration	View the amount of time the root cause analysis (RCA) data is retained in the RCA database.

Table 91: Fields on the Insights Settings Page (Continued)**Fields on the Destination Settings page**

Destination Name	View the name of the destination profile.
------------------	---

Type	View the type of destination selected for a profile.
Attribute	View the email ID or number of reports generated for a destination profile.

Table 91: Fields on the Insights Settings Page (Continued)

Fields on the Scheduler Settings page

Name	View the name of the configured scheduler.
Type	View the type of scheduler selected for a profile.
Start Time	View the date and time to start generating report.
End Time	View the date and time to stop generating reports.
Run For	View the configured duration (in minutes, hours, or days) for which a report is generated.
Repeat	View the frequency of generating a report.

Fields on the Report Settings page

Table 91: Fields on the Insights Settings Page (Continued)

Name	View the name of the generated report.
Format	View the format of the generated report.
Schedules	View when and how often the report is generated.
Destinations	View the location where the report is saved.

Fields on the TSDB Settings page

Table 91: Fields on the Insights Settings Page (Continued)

TSDB Nodes	View the available TSDB nodes in the Paragon Insights installation.
Replication Factor	Enter a value to determine how many copies of the database is needed.
Dedicate	Set the slider to true to dedicate the node to TSDB.
Force	Set the slider to true to ignore all system errors when you remove a failed TSDB node.

Add Alert Blackouts

You can add blackout periods to suppress or mute alarms during scheduled downtimes.

To add blackouts from the Paragon Automation graphical user interface (GUI):

1. Navigate to the **Configuration > Insights Settings** page.
The **Alert Blackout Settings** page is displayed.
2. Click (+) icon to add a new alert blackout setting.
The **Add Alert Blackout** page is displayed.
3. Enter the following information to configure a blackout alert:

Attributes	Description
Blackout Start	Select a blackout start date and time.
Blackout End	Select a blackout end date and time.
Device Group	Select a device group from the drop-down list to apply the blackout configuration.
Attribute	Select an attribute from the drop-down list to apply the blackout configuration.

(Continued)

Attributes	Description
Value	<p>If you have selected an attribute, provide a corresponding associated value. Only the alarms that match this attribute value will be suppressed from the alarms report table.</p> <ul style="list-style-type: none"> • If you have selected Resource as the attribute, select a corresponding value from the Value drop-down list. • If you have selected Resource-Event as the attribute, select a corresponding value from the Value drop-down list. • If you have selected Event as the attribute, enter an event in the Event text box. • If you have selected Topic as the attribute, enter a topic in the Topic text box. <p>For example, for the Resource-Event attribute, you must specify a resource value from the Value drop-down list, as well as specify an Event value. Only those alarms generated by the specified resource that match the Event value will be suppressed from the alarms report table.</p>
Event	<p>Enter the name of the event that you want to apply the blackout configuration to.</p> <p>The field is enabled when you select Event or Resource-Event attributes from the Attribute drop-down list.</p>

4. Click **Save** to save the configuration.

RELATED DOCUMENTATION

[About the Insights Settings Page | 613](#)

[Configure a Notification Profile | 621](#)

[Configure a Retention Policy | 632](#)

[Configure Destination Settings | 633](#)

[Configure Scheduler Settings | 630](#)

[Configure Report Settings | 628](#)

[Manage Time Series Database Settings | 638](#)

About Alert Notifications

Paragon Automation generates alerts that indicate when specific KPI events occur on your devices. To receive alert notifications for these KPI events, you must first configure a notification profile. Once configured, you can enable alert notifications for specific device groups and network groups.

Paragon Automation supports the following notification delivery methods:

- Web Hook
- Slack
- Kafka Publish
- AMQP
- Microsoft Teams
- Email

You can perform the following tasks to enable a device group or network group to send alarm notifications:

- To configure a notification profile. See ["Configure a Notification Profile" on page 621](#).
- To enable alert notifications in device groups or network groups. See ["Enable Alert Notifications for Device Groups and Network Groups" on page 627](#).

Use Exim4 for E-mails

Exim4 is a mail transfer agent (MTA) for Unix-like systems that connect to the internet. The Exim4 agent, that is included in the Paragon Automation software, sends network health reports and alert notifications (for network or device issues) to the e-mail account of the Exim4 host user. The Exim4 host is the Paragon Automation primary node.

NOTE: In case of multinode Paragon Automation installation with more than one primary node, the Exim4 host is one of the primary nodes.

To enable Paragon Automation to use Exim4 MTA, you must do the following:

- Configure the Exim4 hostname (the primary node hostname) in the file that has the environment variables of all microservices. This configuration ensures that Paragon Automation can reach the

Exim4 host when an alert or a report is generated. If the Exim4 host is not reachable, Paragon Automation does not forward the e-mail to the Exim4 agent.

- Configure your DNS server to resolve the Exim4 host's FQDN to the virtual IP address (VIP) of the Paragon Automation Ingress Controller. The format of the FQDN is `hostname.domain.top-level-domain`. This configuration ensures that the Exim4 agent discovers the DNS mail server for the domain, based on the Exim4 host's FQDN. The Exim4 agent then forwards the e-mail to the DNS mail server.

RELATED DOCUMENTATION

| [Configure the Exim4 Agent to Send E-mail](#) | 620

Configure the Exim4 Agent to Send E-mail

To configure the Exim4 Agent to send email, you must first configure the Paragon Automation primary node as the Exim4 host. To do so, you must enter the node's FQDN in the `healthbot.sys` file. The `healthbot.sys` file contains a list of all environment variables for the Paragon Automation microservices. After you modify the `healthbot.sys` file with the Exim4 hostname, you must run the `healthbot restart` command to restart the `alerta` microservice with the changes you made.

NOTE: Before you begin, configure your DNS server to resolve the FQDN of the Exim4 host to the Ingress Controller's VIP.

Use the following steps to enable the Exim4 agent to send e-mails.

1. Log into the Paragon Automation primary node using your server credentials.
2. Type the following command to become a root user.


```
root@primary-node# su -root
```
3. Change the path to `/var/local/healthbot`.


```
root@primary-node# cd /var/local/healthbot/
```
4. Type the following command to open the `healthbot.sys` file.


```
root@primary-node:/var/local/healthbot# vi healthbot.sys
```
5. Scroll down to find the `api-server` section in the `healthbot.sys` file.
6. Type the Paragon Automation primary node FQDN as value for the `HOST_HOSTNAME` variable. For example, `HOST_HOSTNAME = example.domain.top-level-domain`.

7. Type `:wq!` to save the changes and exit the file.
8. Type the following command to update and restart the alerta microservice.

```
root@primary-node:/var/local/healthbot# ./healthbot restart --device-group healthbot -s alerta
```

Now, you have enabled the Exim4 agent to send e-mails to the e-mail account associated with the primary node.

To send e-mail alert notifications, you must configure your e-mail address in a notification profile and enable that notification profile on device groups. For more information, see "[Configure a Notification Profile](#)" on page 621.

To e-mail reports, you must configure your e-mail address in the report settings. For more information, see "[Configure Destination Settings](#)" on page 633.

RELATED DOCUMENTATION

| [Use Exim4 for E-mails](#) | 619

Configure a Notification Profile

Configure a notification profile to define the delivery method used for sending notifications.

To configure a notification profile from the Paragon Automation graphical user interface (GUI):

1. Click **Configuration > Insights Settings**.
The **Insights Settings** page is displayed.
2. Click the **Notification** tab to view the **Notification Settings** page.
3. Click (+) icon to add a new notification.
The **Add Notification Setting** page is displayed.
4. Enter the following information in the **Add Notification Setting** page:

Table 92: Configure Notification Profile

Attributes	Description
Name	Enter a name for the notification.
Description	Enter a description for the notification.

Table 92: Configure Notification Profile *(Continued)*

Attributes	Description
Notification Type	Select a notification type: <ul style="list-style-type: none"> • Web Hook • Slack • Kafka Publish • AMQP • Microsoft Teams • EMails

The notification attributes vary based on the notification type selected.

- If you have selected **Web Hook** notification type, enter the following information:
 - **URL** Enter the URL where the Web Hook notification must be posted.
 - **Username** Enter the username for basic HTTP authentication.
 - **Password** Enter the password for basic HTTP authentication.
- If you have selected **Slack** notification type, enter the following information:
 - **URL** Enter the URL where the Slack notification must be posted.

 This is different from your Slack workspace URL. To create a Slack API endpoint URL, click <https://slack.com/services/new/incoming-webhook> and sign in to your Slack workspace.
 - **Channel** Enter the channel on which the notification must be posted.
- If you have selected **Kafka Publish** notification type, enter the following information:

- **Bootstrap Servers** Enter Kafka host:port pairs to establish the initial connection to the Kafka cluster.
- **Use Hash Partitioner** Select this check box to store messages in consistent partitions based on a partitioning scheme (key).
- **Topic** Enter the name of the Kafka topic in which data must be published.

By default, the Kafka topic naming convention for device group alarm notifications is *device-group.device-id.topic.rule.trigger*.

- Depending on the authentication protocols being used, the required authentication parameters are listed here:

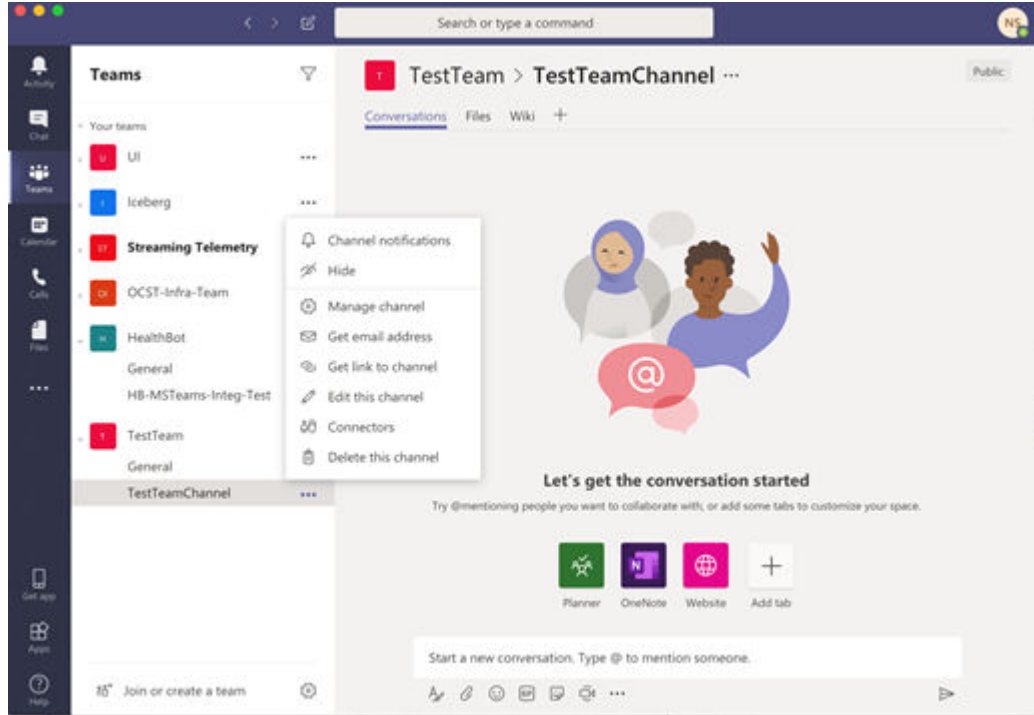
Protocol	Required Parameters
SASL/SSL	Username, password and certificate
SASL/Plaintext	Username and password
SSL	Certificate
Plaintext	None

- **Username** Enter username for SASL/SSL or SASL/plaintext authentication.
 - **Password** Enter password for SASL/SSL or SASL/plaintext authentication.
 - **Certificate** Select Kafka server's CA certificate from the drop-down list.
 - **Upload Certificate** Click **Choose files** and navigate to the location of the file that you want to upload. The file must be in Privacy Enhanced Mail (.pem) format.
- If you selected **AMQP Publish** notification type, enter the following information:
 - **Host** Enter the hostname or IP address of the AMQP server.


- **Port** Enter the port number of the AMQP server.
- **Exchange** Enter the AMQP exchange name (routing agent name).
- **Virtual Host** Enter the virtual host that you configured in RabbitMQ management portal.
- **Routing Key** Enter the AMQP routing key that helps exchange decide how to route the message.
- **Username** Enter the username for the SASL/SSL authentication.
- **Password** Enter the password for the SASL/SSL authentication.
- **CA Profile** Select the CA certificate of AMQP server.
You can configure the CA profile in **Settings > Security**.
- **Local Certificate** Select the local certificate of the AMQP server.
You can configure the local certificate in **Settings > Security**
- **Server Common Name** Enter the common name used while creating the server certificate.
- If you have selected **Microsoft Teams** notification type, enter the following information:
 - **Channel** Paste the web hook URL generated by Teams.

To generate a web hook URL from Teams:

- a. Select the desired channel and click the ellipsis (...).
- b. Click **Connectors** in the displayed menu.



c. Use the **Incoming Webhook** option and click **Configure**.

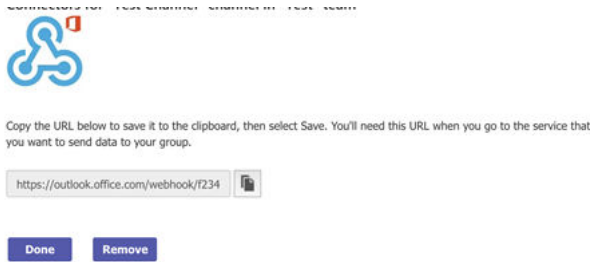


d. Click **Create**.

Customize the image to associate with the data from this Incoming Webhook.



- e. Once the web hook is successfully created, copy the displayed URL.



- If you selected **EMails** notification type, enter the following information:
 - **Email Addresses** Enter an email address and click **Add <email-address>**. You can add one or more email addresses.
 - **Rule Filters** Enter a filter and click **Add <rule-filters>** to define a rule filter. You can add one or more than one rule filter.

Define a rule filter to narrow the scope of what triggers an email.

Examples of rule filters are:

- **interface.statistics/check-interface-flaps**—sends notifications only for the rule check-interface-flaps.
- **system.processes/.*** , **system.cpu/.*** , and **interface.statistics/.*** sends notifications for all rules under the topics system.processes, system.cpu, and interface.statistics.

NOTE: Paragon Insights includes its own mail transfer agent (MTA), so no other mail server is required.

5. Click **Save** to only save the notification settings.

Click **Save & Deploy** to save and deploy the notification.

Apply the notification profile to a device group or network group as shown in ["Enable Alert Notifications for Device Groups and Network Groups" on page 627](#).

RELATED DOCUMENTATION

[Add Alert Blackouts | 617](#)

[About the Insights Settings Page | 613](#)

[Configure a Retention Policy | 632](#)

[Configure Destination Settings | 633](#)

[Configure Scheduler Settings | 630](#)

[Configure Report Settings | 628](#)

[Manage Time Series Database Settings | 638](#)

Enable Alert Notifications for Device Groups and Network Groups

You can enable alert notifications for device groups, network groups, or both by specifying notification profiles, which were previously configured. For more information about alert notifications, see ["About Alert Notifications" on page 619](#).

To enable alert notifications:

- For device groups:

1. Select the **Configuration > Device Group** from the Paragon Automation menu.

The Device Group Configuration page appears.

2. Select the device group and click **Edit** (pencil icon).

The **Edit *Device-Name*** page appears.

3. Navigate to the Notifications section (under Advanced) by clicking the > icon.

4. Depending on the severity of the alerts for which you want to generate notifications, select the notification profiles in the Major, Minor, and Normal fields.

You can select one or more notification profiles in each alert level.

5. Do one of the following:

- **Save** — Save your edits. If you only save configuration changes, you do not deploy the updated configuration. You can use this option when, for example, you are making several changes and want to deploy all your updates at the same time later.

If you choose only to save the changes, you can either commit or roll back the changes later.

For more information, see ["Commit or Roll Back Configuration Changes in Paragon Insights" on page 185](#).

- **Save and Deploy** — Save the configuration and deploy the configuration changes.

The notification profiles selected for each alert severity level generates alerts for the device group. You can view the alerts on the Alerts page (**Monitoring > Alarms and Alerts > Alerts**).

- For network groups:

1. Select the **Configuration > Network** from the Paragon Automation menu.

The Network Configuration page appears.

2. Select the network group and click **Edit** (pencil icon).

The **Edit *Network-Group*** page appears.

3. Navigate to the Notifications section (under Advanced) by clicking the > icon.

You can see the Minor, Major, and Normal fields.

4. Depending on the severity of the alerts for which you want to generate notifications, select the notification profiles in the Major, Minor, and Normal fields.

You can select one or more notification profiles for each alert level.

5. Do one of the following:

- **Save** – Save the changes in configuration. If you only save configuration changes, you do not deploy the updated configuration. You can use this option when, for example, you are making several changes and want to deploy all your updates at the same time later.

If you choose only to save the changes, you can either commit or roll back the changes later.

For more information, see "[Commit or Roll Back Configuration Changes in Paragon Insights](#)" on [page 185](#).

- **Save and Deploy** – Save the configuration changes and deploy the changes.

The notification profiles selected for each alert severity level generates alerts for the network group. You can view the alerts on the Alerts page (**Monitoring > Alarms and Alerts > Alerts**).

RELATED DOCUMENTATION

[About the Alerts Page | 850](#)

[About Alert Notifications | 619](#)

Configure Report Settings

You can generate and download a report on-demand for a device group or a network group by using the Paragon Automation GUI. You can download a report or receive it by email in HTML, JSON, and PDF formats. You must configure report settings before you can generate or download any report.

To configure report settings:

NOTE: Before you configure report settings, ensure that you have configured destination settings from the Destination Settings page and configured a schedule from the Scheduler Settings page.

1. Navigate to the **Configuration > Insights Settings** page, and click the **Report** tab.

The Report Settings page appears.

2. Click the add report (+) icon to add new report settings.

The Add a Report Setting page appears.

3. Enter the following information:

Table 93: Report Setting Information

Attributes	Description
Name	Enter a name for the report.
Format	Select the format in which you want to receive the report. Options: HTML, JSON, PDF.
Schedule(s)	Select a schedule from the drop-down list. To create a schedule for a report, see "Configure Scheduler Settings" on page 630.
Destination(s)	Select a destination profile from the drop-down list. To configure destination settings for a report, see "Configure Destination Settings" on page 633.
Canvas(es)	Select graph canvases to include in the report. Based on the canvas selected, the list of graph panels in the Panel(s) drop-down list changes.
Panel(s)	Select the desired graph panels to include in the report. NOTE: JSON reports include raw time series data only and no graphs.

Table 93: Report Setting Information (Continued)

Attributes	Description
Add Fields	Select the device group and device from the drop-down lists. You can also select rules and topics associated with devices from the subsequent drop-down lists.

4. Do one of the following:

- Click **Save** to only save the configuration to the database without applying changes.

You must commit (or rollback) the configuration changes later. For more information, see "[Commit or Roll Back Configuration Changes in Paragon Insights](#)" on page 185

- Click **Save and Deploy** to save configuration changes to the database and to apply the changes.

RELATED DOCUMENTATION

[Configure Destination Settings | 633](#)

[Configure Scheduler Settings | 630](#)

[About the Insights Settings Page | 613](#)

Configure Scheduler Settings

You can set the date, time, duration, type, and number of times a report is generated from the Scheduler Settings page.

To configure scheduler settings:

1. Navigate to the **Configuration > Insights Settings** page, and click the **Scheduler** tab.

The Scheduler Settings page appears.

2. Click the add schedule (+) icon to add a new schedule.

The Add Scheduler page appears.

3. Enter the following information:

Table 94: Configure Scheduler Settings

Attributes	Description
Name	Enter a name for the schedule.
Scheduler Type	Select the type of schedule from the drop-down list. Options: continuous , discrete .
Run For	Set the duration (in minutes, hours, or days) for which a report is generated. This field is enabled when you select discrete Scheduler Type.
Start On	Set the date and time to start generating reports.
End On	Set the date and time to stop generating reports. To generate reports indefinitely, leave this field blank.
Repeat	Select any one of the following options: <ul style="list-style-type: none"> • Select Never to generate the report only once. • Select Every day to generate daily reports. • Select Every week to generate weekly reports. • Select Every month to generate reports every month. • Select Every year to generate yearly reports. • Select Custom and use the Repeat Every text boxes to configure a custom frequency.

4. Do one of the following:

- Click **Save** to only save the configuration to the database without applying changes.

You must commit (or rollback) the configuration changes later. For more information, see "[Commit or Roll Back Configuration Changes in Paragon Insights](#)" on page 185

- Click **Save and Deploy** to save configuration changes to the database and to apply the changes.

RELATED DOCUMENTATION

[Configure Destination Settings | 633](#)

[About the Insights Settings Page | 613](#)

[Configure Report Settings | 628](#)

Configure a Retention Policy

You can use the Paragon Automation graphical user interface (GUI) to configure a retention policy for time series data that is used for root cause analysis (RCA).

Retention policy is the amount of time the root cause analysis (RCA) data is retained in the RCA database. By default, data is retained for seven days.

The following default retention policies are available in Paragon Automation:

- **CTRL_DAILY_ROLLUP_RETENTION**—Defines for how long the daily rollup data can be stored; default is 1000 days.
- **CTRL_HOURLY_ROLLUP_RETENTION**—Defines for how long the hourly rollup data can be stored; default is 180 days.
- **CTRL_RAW_RETENTION**—Defines for how long the raw data can be stored; default is 14 days.

To configure a retention policy:

1. Click **Configuration > Insights Settings**.
The **Insights Settings** page is displayed.
2. Click the **Retention Policy** tab to view the **Retention Policy Settings** page.
3. Click (+) icon to add a retention policy.
4. Enter the following information:

Table 95: Configure Retention Policy

Attributes	Description
Name	Enter a name for the retention policy.
Duration	Enter the duration, in hours or days, for the retention policy. For example, 1 day is entered as 1d or 24h.

- Click **Save** to only save the configuration.
Click **Save and Deploy** to save and deploy the configuration immediately.
You can now apply the retention policy to a device group.

RELATED DOCUMENTATION

[Add Alert Blackouts | 617](#)

[Configure a Notification Profile | 621](#)

[About the Insights Settings Page | 613](#)

[Configure Destination Settings | 633](#)

[Configure Scheduler Settings | 630](#)

[Configure Report Settings | 628](#)

[Manage Time Series Database Settings | 638](#)

[Add a Data Rollup Summarization Profile | 653](#)

[Add a Raw Data Summarization Profile | 651](#)

Configure Destination Settings

Use the Destination Settings page to configure where you want to send a report. You can send an email of the report, and also save a copy of the report on the server.

To configure destination settings:

- Navigate to the **Configuration > Insights Settings** page, and click the **Destination** tab.
The Destination Settings page appears.
- Click add destination setting (+) icon to configure new destination settings.
The Add Destination page appears.
- Enter the following information:

Table 96: Configure Destination Settings

Attributes	Description
Destination Name	Enter a name for the destination setting. The destination setting name cannot be changed after the settings have been deployed.

Table 96: Configure Destination Settings (Continued)

Attributes	Description
Destination Type	Select the destination type. Options: email , disk .
Email	Enter the email address to which the report must be sent. This text box is displayed when you select Email as the destination type. NOTE: Using the email option also saves a copy of the report to disk.
Max Reports	Specify how many versions of the report is stored on the server. Default value: 5 Older reports are deleted as newer reports are generated and saved. This text box is displayed when you select disk as the destination type.

4. Do one of the following:

- Click **Save** to only save the configuration to the database without applying changes.

You must commit (or rollback) the configuration changes later. For more information, see "[Commit or Roll Back Configuration Changes in Paragon Insights](#)" on page 185

- Click **Save and Deploy** to save configuration changes to the database and to apply the changes.

RELATED DOCUMENTATION

[About the Insights Settings Page | 613](#)

[Configure Scheduler Settings | 630](#)

[Configure Report Settings | 628](#)

Time Series Database (TSDB) Overview

IN THIS SECTION

- [Paragon Insights Microservice | 635](#)
- [TSDB Elements | 635](#)

Paragon Insights collects a lot of time-sensitive data through its various ingest methods. This is why Paragon Insights uses a time-series database (TSDB) to store and manage all of the information received from the various network devices. This topic provides an overview of the TSDB.

Paragon Insights Microservice

Paragon Insights uses Kubernetes for clustering its docker-based microservices across multiple physical or virtual servers (nodes). Kubernetes clusters consist of a primary node and multiple worker nodes. During the Healthbot setup portion of Paragon Insights multinode installations, the installer asks for the IP addresses (or hostnames) of the Kubernetes primary node and worker nodes. You can add as many worker nodes to your setup as you need. However, the number of nodes you add must be more than the value of the replication factor.

TSDB Elements

Paragon Insights supports the following TSDB elements to provide TSDB high availability (HA).

Database Sharding

Database sharding refers to selectively storing data on certain nodes. This method distributes the data among available TSDB nodes and permits greater scaling. This ensures that a TSDB instance handles only a portion of the time series data from the devices.

To achieve sharding, Paragon Insights creates one database per device group/device pair and writes the resulting database to a system determined instance of TSDB hosted on one (or more) of the Paragon Insights nodes.

For example, consider that we have two devices, D1 and D2, and two device groups, G1 and G2. If D1 resides in groups G1 and G2, and D2 resides only in group G2, then we end up with 3 databases: G1:D1, G2:D1, and G2:D2. Each database is stored on its own TSDB instance on a separate Paragon Insights node as shown in [Figure 47 on page 636](#). When a new device is onboarded and placed within a device group, Paragon Insights chooses a TSDB database instance on which to store that device data.

Figure 47: Distributed TSDB




Figure 47 on page 636, shows 3 Paragon Insights nodes. Each of these nodes have a TSDB instance and other Paragon Insights services running.

NOTE:

- A maximum of 1 TSDB instance is allowed on any given Paragon Insights node. Therefore, a Paragon Insights node can have 0 or 1 TSDB instances at any time.
- A Paragon Insights node can be dedicated to running only TSDB functions. No other Paragon Insights functions can run on nodes dedicated to running TSDB functions. This prevents other Paragon Insights functions from starving the TSDB instance of resources.
- We recommend that you dedicate nodes to TSDB to provide the best performance.
- Paragon Insights and TSDB nodes can be added to a running system using the Paragon Insights CLI.

Database Replication

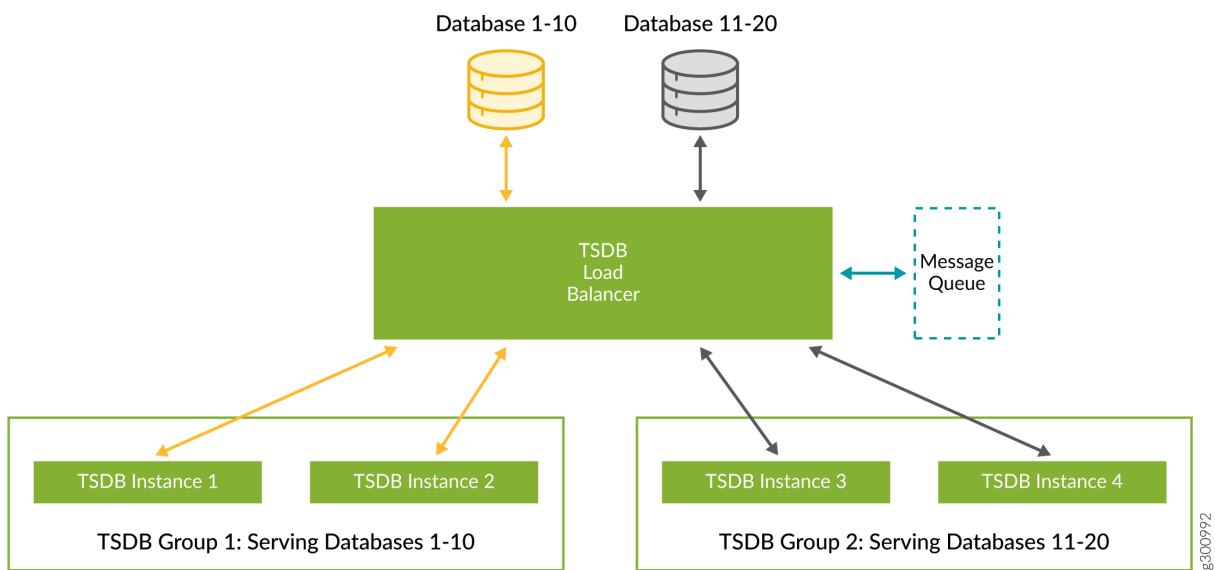
As with any other database system, replication refers to storing the data in multiple instances on multiple nodes. In Paragon Insights, we configure a replication factor to determine how many copies of the database are needed.

A replication factor of 1 creates only one copy of data, and therefore, provides no HA. When multiple Paragon Insights nodes are available and replication factor is set to 1, then only sharding is achieved.

The replication factor determines the minimum number of Paragon Insights nodes needed. A replication factor of 3 creates three copies of data, requires at least 3 Paragon Insights nodes, and provides HA. The higher the replication factor, the stronger the HA and higher the resource requirements in terms of Paragon Insights nodes. If you want to scale your system further, you must add Paragon Insights nodes in exact multiples of the replication factor. For example, 3, 6, 9, etc.

Consider an example where, based on device/device-group pairing mentioned earlier, Paragon Insights has created 20 databases. The Paragon Insights system in question has a replication factor of 2 and has 4 nodes running TSDB. Based on this, two TSDB replication groups are created; in our example they are *TSDB Group 1* and *TSDB Group 2*. In [Figure 48 on page 637](#), the data from databases 1-10 is being written to TSDB instances 1 and 2 in TSDB group 1. Data from databases 11-20 is written to TSDB instances 3 and 4 in TSDB group 2. The outline around the TSDB instances represents a TSDB replication group. The size of the replication group is determined by the replication factor.

Figure 48: TSDB Databases



Database Reads and Writes

As shown in [Figure 47 on page 636](#), Paragon Insights can make use of a distributed messaging queue. In cases of performance problems or errors within a given TSDB instance, this allows for writes to the database to be performed in a sequential manner ensuring that all data is written in proper time sequence.

All Paragon Insights microservices use standardized database query (read) and write functions. This can be used even if the underlying database system is changed at some point in the future. This allows for flexibility in growth and future changes. Other read and write features of the database system include:

- In normal operation, database writes are sent to all TSDB instances within a TSDB group.
- Database writes can be buffered up to 1GB per TSDB instance so that failed writes can be retried until successful.
- If problems persist and the buffer fills up, the oldest data is dropped in favor of new data.
- When buffering is active, database writes are performed sequentially so that new data cannot be written until the previous write attempts are successful.
- Database queries (reads) are sent to the TSDB instance which has reported the fewest write errors in the last 5 minutes. If all instances are performing equally, then the query is sent to a random TSDB instance in the required group.

RELATED DOCUMENTATION

[Manage Time Series Database Settings | 638](#)

[Backup and Restore the TSDB | 640](#)

Manage Time Series Database Settings

You can use the Paragon Automation GUI to configure the time series database (TSDB) settings.

To configure TSDB settings:



WARNING: Selecting, deleting, or dedicating TSDB nodes must be done during a maintenance window because some services will be restarted and the Paragon Automation GUI will likely be unresponsive.

1. Select **Configuration** > **Insights Settings**.

The **Insights Settings** page appears.

2. Click **Time Series Database**.

The **TSDB Settings** tabbed page appears.

3. From the **TSDB Settings** tabbed page, you can:

- a. Select one or more nodes (from the **TSDB Nodes** list) to be used as TSDB nodes.

(The **TSDB Nodes** list displays the available nodes in the Paragon Automation installation that you can select as TSDB nodes. By default, Paragon Automation automatically selects one node as a TSDB node.)

- b. Set the replication factor by typing a value (or by using the arrows to specify a value) in the **Replication Factor** text box.

(The replication factor determines how many copies of the database are needed. The replication factor is set to 1 by default.)

- c. Dedicate nodes as TSDB nodes by clicking the **Dedicate** toggle to turn it on.

A TSDB node might have more than one microservice running. However, when you dedicate a node as TSDB node, it runs only the TSDB microservice, and stops running all other microservices.

NOTE:

- If the node is associated to a persistent volume (storage in a cluster), then you cannot use that node as a dedicated TSDB node.
- A fail-safe mechanism ensures that you cannot dedicate all Paragon Automation nodes as as TSDB nodes.

- d. Ignore system errors (when you remove or replace a failed TSDB node from Paragon Automation) by clicking the **Force** toggle to turn it on.

For example, when a TSDB node fails and the replication factor for that node is set to one, the TSDB data for that node is lost. In this scenario, the failed TSDB node must be removed from Paragon Automation. However, when you try to replace the failed node with a new node, the backup of the node fails with a system error because the replication factor was set to one. If you want to proceed with replacing the node, you must turn the **Force** toggle on.

- e. Delete a node that was previously assigned as a TSDB node by clicking **X** next to the name of the TSDB node.

The node is removed as a TSDB node when you deploy the new configuration changes.

4. Do one of the following:

- Click **Save** to only save the configuration changes to the database without applying the changes to the TSDB nodes.

You must commit (or rollback) the configuration changes later. For more information, see "[Commit or Roll Back Configuration Changes in Paragon Insights](#)" on page 185.

- Click **Save & Deploy** to save configuration changes to the database and to apply the changes to the TSDB nodes.

5. In the pop-up that appears, click **OK** to confirm.

You are returned to the **TSDB Settings** tabbed page.

RELATED DOCUMENTATION

[Time Series Database \(TSDB\) Overview | 635](#)

[About the Insights Settings Page | 613](#)

Backup and Restore the TSDB

You can backup and restore the Time Series Database (TSDB) separately from other configuration elements. You must set a user-defined `HB_EXTRA_MOUNT1` environment variable on the Paragon Insights server prior to any back up or restore operation.

For example, use the following command to set the environment variable.

```
export HB_EXTRA_MOUNT1=/root/.kube/config
```

Here, `HB_EXTRA_MOUNT1` is a variable and is user defined.

The backup and restore operations for the TSDB are available only through the Paragon Insights CLI. The backup and restore commands are invoked by using a predefined python script, `healthbot.py`. You must have root access to the CLI interface of the Paragon Insights server in order to issue these commands.

The generic command along with the required and optional arguments (in square brackets), for performing backup and restore, is described here with examples.

```
/var/local/healthbot/healthbot tsdb (backup|restore) [-h] [--database DATABASE] [--all] --path PATH
```

The required arguments for the `/var/local/healthbot/healthbot tsdb` command are:

- `backup`—perform a backup operation
- `restore`—perform a restore operation
- `--path`
 - For backup: `PATH_TO_DIR`—path to the directory where the backup file will be generated.
 - For restore: `PATH_TO_DIR/BACKUP_FILENAME`—absolute path of backup file that needs to be restored.
- `/var/local/healthbot/healthbot tsdb stop-services`—stop service
- `/var/local/healthbot/healthbot tsdb start-services --port portnumber`—start service

NOTE: In Paragon Automation, the TSDB port is not exposed by default. External API queries to TSDB do not need this port to be exposed as well. However, if you use external tools such as Grafana, or you need run a query to the TSDB directly (and not through APIs), the TSDB port must be exposed by using the `/var/local/healthbot/healthbot tsdb start-services` command.

The optional arguments for the `/var/local/healthbot/healthbot tsdb` command are:

```
-h, --help          show this help message and exit
--database DATABASE Takes backup (or restore) of the given list of databases. Either
                    database or all flag must be configured
--all              Takes backup (or restores) of all the databases. Either database or
                    all flag must be configured
```

Back up the TSDB

```
/var/local/healthbot/healthbot tsdb backup --all --path PATH_TO_DIR
```

Example – Back up the TSDB and Store it in *HB_EXTRA_MOUNT3*

```
/var/local/healthbot/healthbot tsdb backup --all --path $HB_EXTRA_MOUNT3
```

Restore the TSDB

```
/var/local/healthbot/healthbot tsdb restore --all --path PATH_TO_DIR/BACKUP_FILENAME
```

Example – Restore the TSDB from *HB_EXTRA_MOUNT2*

```
/var/local/healthbot/healthbot tsdb restore --all --path $HB_EXTRA_MOUNT2
```

RELATED DOCUMENTATION

| [Time Series Database \(TSDB\) Overview](#) | 635

Time Series Database Replication Scenarios

SUMMARY

IN THIS SECTION

- [Points to Remember | 642](#)
- [Scenario One | 643](#)
- [Scenario Two | 644](#)
- [Scenario Three | 645](#)
- [Frequently Asked Questions | 645](#)

Paragon Insights collects a lot of time-sensitive data through its various ingest methods. Paragon Insights uses a time series database (TSDB) to store and to manage this information received from the various network devices. For more information about TSDB and on managing TSDB settings, see ["Manage Time Series Database Settings" on page 638](#).

These topics explain the various scenarios that you might come across after you have configured TSDB settings from the Paragon Automation GUI.

Points to Remember

1. TSDB Nodes

- The **TSDB Nodes** list in the **TSDB Settings** page of the GUI displays the available nodes in the Paragon Automation installation that you can select as TSDB nodes. By default, Paragon Automation automatically selects one node as a TSDB node.
- A TSDB node might have more than one microservice running. However, when you dedicate a node as TSDB node, it only runs the TSDB microservice, and stops running all other microservices.
- If the node is associate with a persistent volume (storage in a cluster), then you cannot use that node as a dedicated TSBD node.
- A fail-safe mechanism ensures that you cannot dedicate all Paragon Automation nodes as TSDB nodes.
- When a TSDB node fails, you can rebuild the damaged server or component. However, if the replication factor is set to one, the TSDB data for the node is lost.
- You can ignore system errors when you remove or replace a failed TSDB node from Paragon Automation.

- Selecting, deleting, or dedicating TSDB nodes must be done during a maintenance window because some services will be restarted and the Paragon Automation GUI will likely be unresponsive.

2. Replication Factor

- Replication refers to storing data on multiple instances on multiple nodes. In Paragon Automation, we configure a replication factor to determine how many copies of the database are needed. The replication factor determines the minimum number of Paragon Automation nodes needed.
- The replication factor is set to 1 by default. A replication factor of 1 creates only one copy of data, and therefore, provides no high availability (HA). A replication factor of 3 creates three copies of data, requires at least 3 Paragon Automation nodes, and provides HA.
- The higher the replication factor, the stronger the HA and higher the resource requirements in terms of Paragon Automation nodes.
- If you want to scale your system further, you must add Paragon Automation nodes in multiples of the replication factor.

For example, when you set the replication factor to three, you can add Paragon Automation nodes in multiples of three, such as three, six, nine, etc.

3. Database Sharding

- Database sharding refers to selectively storing data on certain nodes. This method distributes the data among available TSDB nodes and improves scaling.

4. Database Reads

- Database queries (reads) are sent to the TSDB instance which has reported the fewest write errors in the last 5 minutes. If all instances are performing equally, then the query is sent to a random TSDB instance in the required group.

For more information, see ["Time Series Database \(TSDB\) Overview" on page 635](#).

Scenario One

Consider the following TSDB configuration:

Table 97: TSDB Replication Scenario One

Number of databases	20
Replication factor	2

<p>TSDB nodes</p>	<p>4 (TSDB-1, TSDB-2, TSDB-3, TSDB-4)</p> <p>You can specify TSDB nodes from the TSDB Settings page. The replication factor is set to one by default. When the replication factor is set to one, Paragon Automation will select one TSDB node to store a copy of the data of a database.</p> <p>However, in this scenario, when the replication factor is set to two, Paragon Automation will store a copy of the data on two different TSDB nodes.</p>
<p>TSDB groups <i>(created automatically by Paragon Automation)</i></p>	<p>2 TSDB groups with 2 TSDB nodes each</p> <p>Two TSDB groups are created automatically by Paragon Automation after you set the replication factor and add TSDB nodes.</p>

In this scenario, when a new database is created, a TSDB group that serves the least number of databases is automatically assigned to that database. Data from the database is stored (TSDB writes) in this assigned TSDB group. A copy of the data is maintained in both TSDB nodes that form the TSDB group.

Scenario Two

Consider the following TSDB configuration

Table 98: TSDB Replication Scenario Two

<p>Number of databases</p>	<p>20</p>
<p>Replication factor</p>	<p>2</p>
<p>TSDB nodes</p>	<p>8 (TSDB-1, TSDB-2... TSDB-8)</p> <p>You can specify TSDB instances from the TSDB Settings page. The replication factor is set to 1 by default. When the replication factor is set to one, Paragon Automation will select one TSDB node to store a copy of the data of a database.</p> <p>However, in this scenario, when the replication factor is set to two, Paragon Automation will store a copy of the data on two different TSDB nodes.</p>

TSDB groups <i>(created automatically by Paragon Automation)</i>	4 TSDB groups with 2 TSDB nodes each Four TSDB groups are created automatically by Paragon Automation after you set the replication factor and add TSDB nodes.
--	--

In this scenario where the total number of databases are 20, the TSDB group that serves the least number of databases is selected. However, if all TSDB groups serve the same number of databases, the TSDB group is picked at random. A copy of the data is maintained in all TSDB nodes that form the TSDB group.

Scenario Three

Consider the following TSDB configuration

Table 99: TSDB Replication Scenario Three

Number of databases	20
Replication factor	1
TSDB nodes	4 (TSDB-1, TSDB-2... TSDB-4) You can specify TSDB nodes from the TSDB Settings page. In this scenario, the replication factor is set to the default value of one. When the replication factor is set to one, Paragon Automation will select one TSDB node to store a copy of the data of a database.
TSDB groups <i>(created automatically by Paragon Automation)</i>	4 TSDB groups with 1 TSDB node in each group Four TSDB groups are created automatically by Paragon Automation after you set the replication factor and add TSDB nodes.

In this scenario, four TSDB groups are created automatically with one node in each group. Since the replication factor is one, only one copy of data is maintained per TSDB node, and therefore, provides no high availability (HA).

Frequently Asked Questions

1. What happens when one node has gone down and the replication factor is set to one?

Since the replication factor is set to one, only one copy of data is maintained within a TSDB group. When the TSDB node fails, you cannot recover the data. However, the last 1GB of data directed to the failed node will buffer continuously till the TSDB node is replaced or recovered.

2. *What happens when one node has gone down and the replication factor is more than one?*

The data will be stored on other TSDB nodes in the TSDB group. Data can be recovered when you replace the failed TSDB node with a new TSDB node. The last 1GB of data directed to the failed node will buffer continuously till the TSDB node is replaced or recovered.

3. *What do you do when a failed node cannot be recovered?*

When a failed node cannot be recovered, you can replace the failed node with a new node. The removal of a failed node and adding of a new node must be done within the same commit configuration. All data from other TSDB nodes can be copied to the new node if the replication factor is set to more than one.

4. *What do you do when a new node cannot be added?*

When a new node cannot be added, remove the failed node and adjust the replication factor to match new number of nodes. Ensure that the changes are done within the same commit configuration. For example, when the replication factor is set to two and there are four TSDB nodes, you can configure the replication factor to either one or three after you remove the failed node. When you configure the replication factor to one, there will be no HA but there will be sharding. When you configure the replication factor to three, there will be HA but no sharding.

5. *What do you do when there is a disk failure?*

You must replace the failed disk with a new disk. You then remove the failed node and add a new node with the same replication. This will ensure that the disk in the new node will have all the data as the other nodes in the TSDB group. After you bring up the node with the new disk, data is automatically restored.

Data Summarization Profiles

IN THIS CHAPTER

- [Data Summarization Overview | 647](#)
- [About the Raw Data Summarization Profiles Page | 649](#)
- [About the Data Roll Up Summarization Profiles Page | 650](#)
- [Add a Raw Data Summarization Profile | 651](#)
- [Add a Data Rollup Summarization Profile | 653](#)
- [Apply Data Summarization Profiles | 655](#)

Data Summarization Overview

Paragon Insights provide a way to store raw data using summarization profiles to reduce the disk space and improve performance of time series database (TSDB).

Paragon Insights collects data from devices by using push or pull data collection ingest methods. You can create rules or use the available pre-defined rules to determine how and when data is collected. The telemetry data can be summarized as a function of time or when a change occurs.

For time-based data summarization, the raw data points are grouped together into user-defined time spans, and each group of data points is summarized into one data point using aggregate functions.

Paragon Insights also supports data rollup summarization. Data rollup summarization helps you to summarize field-level data. Field-level data is processed data that provides information on network devices and its components, and is stored in fields in the TSDB. A field is a single piece of information that forms a record in a database. In TSDB, multiple fields of processed data make a record. Data rollup summarization enables efficient data storage and also ensures retaining of data for a longer duration.

[Table 100 on page 648](#) provides a list of the supported data summarization algorithms and a description of their output:

Table 100: Descriptions of the Data Summarization Algorithms

Algorithm	Description of output
Latest	Value of the last data point collected within the time span.
Count	Total number of data points collected within the time span.
Mean	Average value of the data points collected within the time span.
Min	Minimum value of the data points collected within the time span.
Max	Maximum value of the data points within the time span.
On-change	Value of the data point whenever the value is different from the previous data point (occurs independently from the user-defined time span).
Stddev	Standard deviation of the data points collected within the time span.
Sum	Sum of the data points collected within the time span.

If no summarization algorithm is associated with the data, the following algorithms are used by default:

Data type	Data summarization algorithm
Float, integer, unsigned	Mean
Boolean, string	On-change

You can use data summarization profiles to apply specific summarization algorithms to raw data and field-level data collected by Paragon Insights for a specific device group:

These topics provide instructions on how to create a data summarization profile.

- [Creating a Raw Data Summarization Profile](#). See "[Add a Raw Data Summarization Profile](#)" on page 651.

- Apply a Data Summarization Profile to a Device Group. See ["Apply Data Summarization Profiles"](#) on page 655.
- Create a Data Rollup Summarization Profile. See ["Add a Data Rollup Summarization Profile"](#) on page 653.

RELATED DOCUMENTATION

[Backup and Restore the TSDB](#) | 640

[Time Series Database \(TSDB\) Overview](#) | 635

About the Raw Data Summarization Profiles Page

IN THIS SECTION

- [Tasks You Can Perform](#) | 649
- [Field Descriptions](#) | 649

Raw data summarization refers to the process of creating a concise version of time series database's (TSDB) raw data. You create a raw data summarization profile to improve the performance and disk space utilization of the TSDB.

To access the **Raw Data Summarization Profiles** page from the Paragon Automation graphical user interface (GUI), click **Configuration > Summarization Profiles > Raw Data**.

Tasks You Can Perform

- Add a raw data summarization profile. See ["Add a Raw Data Summarization Profile"](#) on page 651.
- Edit a raw data summarization profile.
- Delete a raw data summarization profile.

Field Descriptions

The following table describes the fields on the **Raw Data Summarization Profiles** page:

Field	Description
Data Summarization Name	Displays the name of the raw data summarization profile.
Type Aggregate	Displays the name data type (string, integer, boolean, float) that you apply to the raw data summarization profile.
Path Aggregate	Displays the sensor path name that is associated with the raw data summarization profile.

About the Data Roll Up Summarization Profiles Page

IN THIS SECTION

- [Tasks You Can Perform | 650](#)
- [Field Descriptions | 651](#)

Data rollup summarization refers to the process of creating a concise version of time series database's (TSDB) field data. You create a rollup summarization profile to summarize processed data that is stored in fields in the TSDB.

To access the **Data Roll Up Summarization Profiles** page from the Paragon Automation graphical user interface (GUI), click **Configuration > Summarization Profiles > Data Roll Up**.

Tasks You Can Perform

- Add a data rollup summarization profile. See "[Add a Data Rollup Summarization Profile](#)" on page [653](#).
- Edit a data rollup summarization profile.
- Delete a data rollup summarization profile.

Field Descriptions

The following table describes the fields on the **Data Roll Up Summarization Profiles** page:

Field	Description
Name	Displays the name of the data rollup summarization profile.
Data Rollup Order	Displays the interval in which data is summarized, and how long the data is retained.
Rule(s)	Displays the rule(s) that you have applied to the data rollup summarization profile.

Add a Raw Data Summarization Profile

To add a raw data summarization profile that can be applied to a device group:

1. Click **Configuration > Summarization Profiles > Raw Data** link in the left navigation bar.

The Raw Data Summarization Profiles page is displayed.

2. Click **+** to add a summarization profile.

The Add Data Summarization Profile page appears.

3. In the **Name** field, enter the name of the profile.

4. Click **Add Type Aggregate** to add an aggregate type.

The **Name** and **Function** drop-down lists are displayed.

Follow these steps to select a name data type and associate it with a data summarization algorithm.

The algorithm configured for a specific sensor path name overrides the algorithm configured for the corresponding data type.

- a. Select a name data type from the **Name** drop-down list.

The available name data types to choose from are string, integer, boolean, float, and unsigned integer.

You can also select unsigned integer as a name data type. An unsigned integer is a data type that can contain values from 0 through 4,294,967,295.

- b. After you have selected a name data type, you associate it with a data summarization function.

To associate a name data type with a data summarization function, select a function from the Functions drop-down list.

The available functions to choose from are latest, count, mean, min, max, on-charge, stddev, and sum.

- c. (Optional) To add another aggregate type, click **Add Type Aggregate**, and repeat step **4.a** and step **4.b**.

5. Click **Add Path Aggregate** to add an aggregate path.

The **Name** and **Function** drop-down lists are displayed.

To assign a sensor path name and associate it with a data summarization algorithm:

NOTE: The algorithm configured for a specific sensor path name overrides the algorithm configured for the corresponding data type.

- a. Enter a sensor path name in the Name field.

You can enter a path name for a sensor that is not supported by Paragon Insights . For sensors supported by Paragon Insights , the path name must be entered in the following format:

Sensor	Path Name Format	Example
Open Config	<i>sensor-path</i>	/components/component/name
Native GPB	<i>sensor-name:sensor-path</i>	jnpr_qmon_ext:queue_monitor_element_info.percentage
iAgent	<i>yaml-table-name:sensor-path</i>	REutilizationTable:15_min_cpu_idle
SNMP	<i>snmp-table-name:sensor-path</i>	.1.3.6.1.2.1.2.2:jnxLED1Index ospfNbrTable:ospfNbrIpAddr
Syslog	<i>pattern-set: sensor-path</i>	interface_link_down:operational-status
Flow (NetFlow)	<i>template-name:sensor-path</i>	hb-ipfix-ipv4- template:sourceIpV4Address

- b. After you have entered a sensor path name, you associate it with a data summarization function.

To associate a sensor path name with a data summarization function, select a function from the Functions drop-down list.

The available functions to choose from are latest, count, mean, min, max, on-charge, stddev, and sum.

- c. (Optional) To add another aggregate path, click **Add Path Aggregate**, and repeat step **5.a** and step **5.b**.
6. Click **Save** to save the configuration or click **Save and Deploy** to save and deploy the configuration.
7. You can now apply the raw data summarization profile that you added to a specific device group. For more information, see "[Apply Data Summarization Profiles](#)" on page 655.

Add a Data Rollup Summarization Profile

Data summarization refers to the process of creating a concise version of raw data and field data. Data can be summarized as a function of time or when a change occurs. You can add a rollup summarization profile to summarize processed data that is stored in fields in the TSDB.

You can add a data rollup summarization profile to apply to a device group from the:

- Paragon Automation GUI
- MGD CLI

To add a data rollup summarization profile:

1. Click **Configuration > Summarization Profiles > Data Roll Up** link in the left navigation bar.
The Data Roll Up Summarization Profiles page appears.
2. Click the Add (+) icon.
The Add Data Rollup Summarization Profile page appears.
3. Enter the name of the profile in the **Name** text box.
4. Click **Add Rule** to add a rule for the device profile.
The **Name** and **Apply on Existing Data** lists are displayed.

Follow these steps to select a rule, to apply the rule to the profile, and to apply the rule to existing data.

- a. Select a rule to apply to the profile from the **Name** list.
- b. To apply the rule that you selected to existing data, select **True** from the **Apply on Existing Data** list.
The default value is **False**.

- c. Click **Add Field** to associate a default field to an aggregate function.
To associate a default field to an aggregate function:
 - i. Select a default field from the **Name** list to which you want to apply an aggregate function.
 - ii. Select one or more aggregate functions from the **Aggregate Function** list that you want to apply to a default field.
 - d. (Optional) To add another rule to the profile, click **Add Rule**, and repeat steps **4.a** through step **4.c**.
5. Click **Add Data Rollup Order** to define how multiple rollup orders are configured, retained, and executed.
- The **Name** and **Retention Policy** lists, and the **Rollup Interval** text box are displayed.
- To define a data rollup order:
- a. Enter a name to identify the data rollup order in the **Name** text box.
The maximum length is 64 characters.
Regex pattern: “[a-zA-Z][a-zA-Z0-9_-]*”
 - b. Enter a value in the **Rollup Interval** text box to define an interval in which the data is summarized.
Regex pattern: “[1-9][0-9]*[mhdw]”, where m is minutes, h is hours, d is days, and w is weeks.
Minimum value is 30m. Maximum value is 52w.
 - c. Select the retention policy for the rollup order from the **Retention Policy** list.
A retention policy defines how long you want to retain the new data.

Selecting a retention policy is optional. If you do not select a retention policy, the device group retention policy is considered by default.
 - d. (Optional) To define another data rollup order, click **Add Data Rollup Order**, and repeat steps **5.a** through step **5.c**.
6. Click **Save** to only save the configuration.
Click **Save and Deploy** to save and deploy the configuration immediately.
7. You can now apply the data rollup summarization profile that you add to a specific device group. For more information, see ["Apply Data Summarization Profiles" on page 655](#).

[Figure 49 on page 655](#) is an example configuration of how you can configure a data rollup summarization profile from the CLI.

Figure 49: Example CLI Configuration of Creating a Data Rollup Summarization Profile

```

field-profile sample_profile1 {
  rule rule1 {
    apply-on-existing-data;
    field rule1_field1 {
      aggregate-function max;
    }
    field rule1_field2 {
      aggregate-function [ mean min ];
    }
    field rule1_field3 {
      aggregate-function [ count mean std-dev ];
    }
  }
  rule rule2 {
    field rule2_field1 {
      aggregate-function first;
    }
  }
  rule rule3 {
    apply-on-existing-data;
    field rule3_field1 {
      aggregate-function [ first std-dev ];
    }
  }
  data-rollup-order hourly_rollup {
    interval 1h;
    retention-policy 2weeks;
  }
  data-rollup-order daily_rollup {
    interval 24h;
    retention-policy 2months;
  }
  data-rollup-order weekly_rollup {
    interval 7d;
    retention-policy 6months;
  }
  data-rollup-order monthly_rollup {
    interval 4w;
    retention-policy 1year;
  }
  data-rollup-order yearly_rollup {
    interval 52w;
    retention-policy 5years;
  }
}

```

Apply Data Summarization Profiles

After you create a data summarization profile, you can apply the profile to a specific device group to start summarizing TSDB data:

1. Click **Configuration > Device Groups** in the left-navigation bar.

The **Device Group Configuration** page is displayed.

2. Select the check box next to the name of the device group to which you want to apply the data summarization profile.

3. Click **Edit Device Group** to edit the device group.

The **Edit <device-group-name>** page is displayed.

4. Apply a raw data summarization profile

To apply a raw data summarization profile to a device group:

a. Click **Summarization**.

The **Time Span** and **Data Summarization** text boxes are displayed.

b. Enter the **Time Span** in seconds (s), minutes (m), hours (h), days (d), weeks (w), or years (y).

c. Choose the data summarization profiles from the drop-down list to apply to the ingest data.

To edit or view details about saved data summarization profiles, go to the **Data Summarization** page and click the **Settings** menu in the left-navigation bar.

If you select two or more profiles, the following guidelines apply:

- If the same data type or sensor path name is configured in two or more profiles, the associated algorithms will be combined.
- The table that stores the summarization output includes columns of summarized data for each algorithm associated with each data field collected by Paragon Insights . The naming convention for each column is as follows:

Number of algorithms associated with a data field	Column name for the summarized output
1	<i>field-name</i> Example: 5_sec_cpu_idle
2	<i>field-name_first-algorithm-name, field-name_second-algorithm-name</i> Example: 5_sec_cpu_idle_MIN, 5_sec_cpu_idle_MAX
3	<i>field-name_first-algorithm-name, field-name_second-algorithm-name, field-name_third-algorithm-name..</i>

Apply a data rollup summarization profile

Points to remember before you apply a data rollup summarization profile to a device group:

- Ensure that the rules present in the rollup profile are already associated with the device group.
- You can add one or more rollup summarization profiles to a device group.
- Rules configured across all the profiles associated with a device group must be unique.
- While associating a rollup profile with a device group, the interval of the first data rollup order must be less than the device group retention policy to avoid data overflow. The device group retention policy is set to 7 days by default.
- When you want to remove a rule that is associated to a device group, you must first remove the data rollup summarization profile.

To apply a data rollup summarization profile to a device group:

- Click **Rollup Summarization**.
The **Rollup Summarization Profiles** drop-down list is displayed.
- Select the rollup summarization profiles you want to associate to this device group from the **Rollup Summarization Profiles** drop-down list.
- (Optional) You can also deploy rollup configuration at the device group-level by using the CLI.
See [Figure 50 on page 657](#) for an example CLI configuration.

Figure 50: Example CLI Configuration

```
device-group DG {
    authentication {
        password {
            username root;
            password "$9$0E9lBIhx7VsYohSbs2aiHCtu0IclK8"; ## SECRET-DATA
        }
    }
    field-data {
        rollup {
            profile sample_profile1;
        }
    }
    devices [ d1 d2 ];
}
```

- Click **Save** to only save the configuration.

Click **Save and Deploy** to save and deploy the configuration immediately.

RELATED DOCUMENTATION

| [Data Summarization Overview](#) | 647

8

PART

Configure Your Network

[Topology | 660](#)

[Network Information Table | 688](#)

[Tunnels | 815](#)

[Change Control Management | 824](#)

Topology

IN THIS CHAPTER

- [Interactive Map Features Overview | 660](#)
- [About the Topology Page | 677](#)
- [Left Widget Options on Topology Page | 680](#)
- [Group Nodes | 682](#)
- [Group Nodes and Links into a Topology Group | 683](#)
- [Ungroup Nodes | 684](#)
- [Automatically Group Nodes | 684](#)
- [Manage Map Layouts | 686](#)

Interactive Map Features Overview

IN THIS SECTION

- [Right-click Functions | 661](#)
- [Topology Settings Pane | 669](#)

A topology map is interactive, which means you can use the features within the map to customize the map and the network information table. 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 on the earth.
- **World wrapping/map wrapping:** Scrolling the map in one direction is like spinning a globe. For example, this feature enables the representation of links across an ocean.

Right-click Functions

You can right-click a node, node group, link, or blank space on the topology map to access context-specific menus.

[Table 101 on page 661](#) describes the options that are displayed when you right-click any blank space on the topology map.

Table 101: Right-Click Options for Blank Space on the Topology Page

Option	Description
Grouping You can view, group, or ungroup the nodes in a node group.	
Nodes & Groups	Displays the nodes and node groups in the topology. You can select a particular node group to be displayed on the topology map.
Auto Grouping	Enables you to specify the criteria that are used to automatically group nodes. For more information, see "Automatically Group Nodes" on page 684 .
Group Selected Nodes	Groups the selected nodes. For more information, see "Group Nodes" on page 682 .
Ungroup Selected Nodes	Ungroups the nodes from the selected group. For more information, see "Ungroup Nodes" on page 684 .
Ungroup All	Ungroups the nodes in all groups.
Group Selected Nodes/Links into Topology Group	Groups the nodes or links that are selected on the topology map into a topology group based on the slice ID that you specify.
Clear Topology Group from the Selected Nodes/Links	Removes the slice IDs from the nodes and links that are present within the selected topology group.
Select All Nodes	Selects all the nodes on the topology map. This option is a shortcut to using the Shift-left-click option to create a selection box around all nodes or individually shift-clicking on all nodes.

Table 101: Right-Click Options for Blank Space on the Topology Page *(Continued)*

Option	Description
Show All Nodes and Links	Restores the topology map so that it includes all the nodes and links in the network, as opposed to a filtered subset.
Save to Default Map Layout	Saves the current layout as the default. Setting this option does not change the name of the default in the Manage Layouts page.
Layout	
Manage Map View	Save, load, or edit the current topology map. For more information, see "Manage Map Layouts" on page 686 .
Distribute All Nodes	Select multiple nodes on the topology map and redistribute them to improve visual clarity or for personal preference. Distributes all the nodes in the map, pushing elements away from each other and minimizing overlap.
Distribute Selected nodes	Forces the selected elements away from each other and minimizes overlap.
Circle selected nodes	Arranges the selected nodes in a roughly circular pattern, with the nodes and links separated adequately.
Straighten selected nodes	Aligns the selected nodes in a linear pattern.
Reset by Coordinates	Resets the map to display the nodes based on their configured coordinates (latitude and longitude). NOTE: You can reset the distribution of the nodes on the topology map according to geographical coordinates if you have set the latitude and longitude values of the nodes. It can be useful to have the country map backdrop displayed when you use this distribution model.
Set Coordinates from Map	Reconfigures the node coordinates based on the current location of the nodes on the map.

Table 101: Right-Click Options for Blank Space on the Topology Page (*Continued*)

Option	Description
Import from	Import a layout from a CSV or GeoJSON file. For more information, see https://geojson.org/ .
Export to	Export a layout to a CSV or GeoJSON file.
Node Label	<p>Select one of the following options to label the nodes on the topology map:</p> <ul style="list-style-type: none"> • Name • Hostname • Hostname, SID • Hostname, Slices • IP Address • IPv6 Address • IP, SID • ISIS System ID • OS Version • OSPF ref BW • SID • Type • Slices • Hide Label—Hides all the labels for the network elements on the topology map

Table 101: Right-Click Options for Blank Space on the Topology Page *(Continued)*

Option	Description
Link Label	<p>Select one of the following options to label the links on the topology map:</p> <ul style="list-style-type: none"> • Name • Node Name A::Z • Interface A::Z • IP A::Z • IP, SID A::Z • IPv6, SRv6 SID Function A::Z • TE Metric A::Z • Bandwidth A::Z • Delay A::Z • Interface Util A::Z • ISIS1 Metric A::Z • ISIS2 Metric A::Z • Measured Delay A::Z • OSPF Metric A::Z • Packet Loss Metric A::Z • RSVP Bandwidth A::Z • RSVP Util A::Z • RSVP Live Util A::Z • Shape BW A::Z • SID A::Z • Slices • SRLG • TE Admin Group A::Z • SRv6 SID Function A::Z • Hide Label
Favorites	
Add Selected Nodes to Favorites	Select nodes on the topology map and designate them as favorites.
Remove Selected Nodes from Favorites	Removes the nodes that you select on the topology map from the Favorites tab.
Clear Favorites	Clears all the existing favorite nodes.
Highlight Favorites	Highlights only the favorite nodes on the topology map.

Table 101: Right-Click Options for Blank Space on the Topology Page (Continued)

Option	Description
Show only Favorites	Displays only the favorite nodes on the topology map.
Hide Favorites	Hides all the favorite nodes on the topology map.
Subview	Filters the network elements on the topology map based on Node Type, Autonomous System (AS) number, OSPF Area, ISIS Area, Layer, or Admin Group. For more information about these options, see Table 102 on page 667 .

Table 101: Right-Click Options for Blank Space on the Topology Page *(Continued)*

Option	Description
Timeline	<p data-bbox="589 369 1349 396">Lists activities and status checkpoints, with the most recent activity first.</p> <p data-bbox="589 428 1406 562">You can use the search box at the top of the Timeline drawer to highlight specific events. You can use the up and down arrows to move to the next or previous event (search result). You can click the top arrow (at the bottom-right corner of the drawer) to move back to the topmost event in the timeline.</p> <p data-bbox="589 594 1386 657">You can also refresh the timeline events by clicking the refresh button at the bottom-left corner of the page.</p> <p data-bbox="589 688 1406 823">You can assess the stability of the MPLS network by tracking changes in the number of LSP Up and Down events over time. You can then analyze whether the occurrence of specific other events affects the number of LSP Up and Down events.</p> <p data-bbox="589 854 1162 882">The following event types are included in the Timeline:</p> <p data-bbox="589 913 898 940">Event types related to nodes:</p> <ul data-bbox="589 972 959 1136" style="list-style-type: none"> • PCEP session goes Down • PCEP session goes Up • PCEP session becomes ACTIVE <p data-bbox="589 1167 881 1194">Event types related to links:</p> <ul data-bbox="589 1226 797 1329" style="list-style-type: none"> • Link goes Up • Link goes Down <p data-bbox="589 1360 883 1388">Event types related to LSPs:</p> <ul data-bbox="589 1419 1224 1583" style="list-style-type: none"> • Change in the number of LSPs that are Up • Change in the number of LSPs that are Down • Change in the number of LSPs that are being provisioned <p data-bbox="589 1614 932 1642">Event types related to Controller:</p> <ul data-bbox="589 1673 1049 1776" style="list-style-type: none"> • Path optimization start and end times • Maintenance events start and end times

Table 101: Right-Click Options for Blank Space on the Topology Page (Continued)

Option	Description
Pause/Resume Network Event Processing	<p>Click to pause or resume the processing of network events. When paused, the network information table and topology map are not refreshed in response to network events until you select Resume Network Event Processing. Network events continue to be processed in the background, but they are not refreshed in the UI.</p> <p>This option is beneficial in large networks where the processing of network events results in frequent UI updates.</p>
Reload Network	Reloads the network, and updates the displayed topology map.

Table 102: Subview Options

Options	Description
Node Type	Select the node types from the list.
AS	<p>Assign a color to represent each AS number that is configured on the topology map.</p> <p>From the AS pane, you can select or clear AS numbers by selecting or clearing the corresponding check boxes. Only nodes corresponding to the selected AS numbers are displayed in the topology map.</p>
ISIS Area	<p>Assign a color to represent each IS-IS area identifier that is configured on the topology map. The area identifier is the first three bytes of the ISO network entity title (NET) address.</p> <p>From the ISIS area pane, you can select or clear ISIS area identifiers by selecting or clearing the corresponding check boxes. Only nodes corresponding to the selected area identifiers are displayed in the topology map.</p>
OSPF Area	<p>Assign a color to represent each OSPF area that is configured on the topology map. NONE shows the color assigned to routers that have no OSPF area configured.</p> <p>From the OSPF Areas pane, select or clear OSPF areas by selecting or clearing the corresponding check boxes. Only nodes corresponding to the selected OSPF areas are displayed in the topology map.</p>

Table 102: Subview Options (Continued)

Options	Description
Layer	<p>Include or exclude individual layer information in the topology map.</p> <p>From the Layers list, select the layers (IP, Transport, or both) that you want to display. If you are not using the Multilayer feature, the Layers list contains only IP.</p>
Admin Group	<p>Provides bit-level link coloring options so that you can easily differentiate the different links that are displayed in the topology map.</p> <p>The Admin Group includes manually assigned bit-level attributes that describe the color of the links (up to 32 names or values from bit 0 to bit 32). You can filter by three conditions (all, any, or not). Links with the same color conceptually belong to the same class. You can use this option to implement a variety of policy-based label-switched path (LSP) setups. For more information, see "Assign Names to Admin Group Bits" on page 197.</p>

[Table 103 on page 668](#) describes the right-click options for a node or node group that you select on the topology map.

Table 103: Right-Click Options for Nodes or Node Groups

Option	Description
Filter in Node Table	Filters the nodes that are displayed in the network information table to display the selected nodes or node groups only.
Tunnels On or Thru Node	Opens a new tab in the network information table to show only those tunnels that meet the On or Thru Node criteria.
Tunnels Starting at Node	Opens a new tab in the network information table to show only those tunnels that meet the Starting at Node criteria.
Tunnels Ending at Node	Opens a new tab in the network information table to show only those tunnels that meet the Ending at Node criteria.

NOTE: For the description of other right-click options, see [Table 101 on page 661](#).

[Table 104 on page 669](#) describes the right-click options for a link on the topology map.

Table 104: Right-Click Options for Links

Option	Function
Filter in Link Table	Filters the tunnels that are displayed in the network information table to display only the selected link.
Tunnels On or Thru Link	Opens a new tab in the network information table to show only those tunnels that meet the On or Thru Link criteria.

NOTE: For the description of other right-click options, see [Table 101 on page 661](#).

Topology Settings Pane

You can access the Topology Settings pane by clicking the **Settings** icon in the Topology menu bar that is located at the upper right corner of the Topology Page.

[Table 105 on page 670](#) describes the tabs on the Topology Settings pane.

Table 105: Topology Settings Options

Tab	Description
Nodes	<p>You can perform the following tasks:</p> <ul style="list-style-type: none">• Change the labels of all the nodes on the topology map by selecting an option from the Label list.• After you select a label, click the toggle button to view:<ul style="list-style-type: none">• Background shadow of node labels• Pseudo node label• Only favorite node labels• Isolated nodes• Overload bit marker

Table 105: Topology Settings Options (*Continued*)

Tab	Description
Links	<p>You can perform the following tasks:</p> <ul style="list-style-type: none"> • Change the labels of all the links on the topology map by selecting an option from the Label list. • Click the toggle button to: <ul style="list-style-type: none"> • 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). • Show utilization max instead of Average within bundle: Toggle this button to display the maximum utilization value (in red) instead of the average utilization value for parallel links. For example, if there are two parallel links between two nodes and one link has 10% utilization and the other has 20% utilization, and you want to draw the link as a straight line between the nodes, you can choose which utilization value should be displayed: 15% (average) or 20% (maximum) value. • Show Count of the bundle: Toggle the button to display (in green) how many links are present between the two nodes. • Wrap links as great arcs: Distinguishes links that would have to wrap around the world map. <p>An example is shown in Figure 51 on page 672.</p>

Table 105: Topology Settings Options (*Continued*)

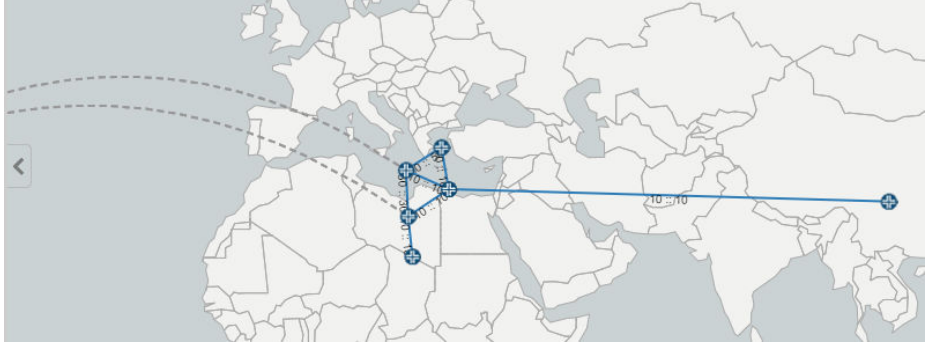
Tab	Description
	<p data-bbox="500 352 1040 386">Figure 51: Wrap Links as Great Arcs: Example</p>  <ul data-bbox="467 827 1414 890" style="list-style-type: none">• Hide Partially Visible links: Hides any link whose end nodes are outside the visible area. This is useful for focusing on a subset of a large network. <p data-bbox="427 926 1398 1024">NOTE: The topology map does not display more than a certain number of node or link labels, even if the topology settings call for labels to be displayed. This constraint improves performance when redrawing a large number of graphic elements.</p>

Table 105: Topology Settings Options *(Continued)*

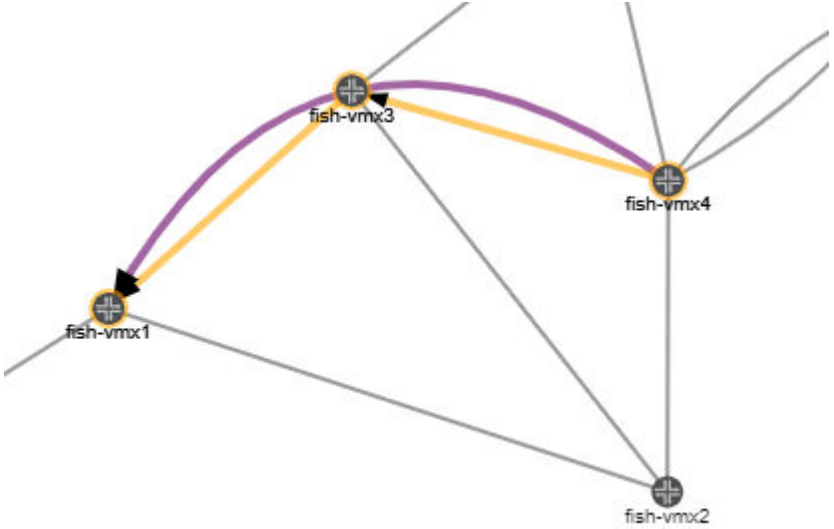
Tab	Description
Tunnels	<p>You can perform the following tasks:</p> <ul style="list-style-type: none"> Click the toggle button to draw a path as a curve, which might improve network visualization. <p>Figure 52 on page 673 shows both curved and straight lines for the same path.</p> <p>Figure 52: Curved and Straight Line Path Depiction</p>  <ul style="list-style-type: none"> Click the toggle button to draw a path through layers if the network includes transport layers.

Table 105: Topology Settings Options (*Continued*)

Tab	Description
General	<p>You can perform the following tasks:</p> <ul style="list-style-type: none">• Click the toggle button to:<ul style="list-style-type: none">• Enable animation while calculating the topology layout.• Show or hide maintenance marker: Displays a red "M" over any link that is part of an active maintenance event.• Show or hide zoom slider: A vertical slider is displayed on the topology menu bar on the top right corner of the topology window.• Enable or disable to zoom to selected node: With this option enabled, when you click the node entry in the network information table (the Node tab), the topology automatically centers the view on that selected node.• Select a Label Size: Select one of the following values as the font size for the node and link labels:<ul style="list-style-type: none">• 8• 10• 12• 14• 16• 18• 20• Apply Opacity effects: Move this slider to select the percent opacity for topology map elements that are not highlighted. Figure 53 on page 675 shows 100% and 20% opacity for comparison.

Table 105: Topology Settings Options (Continued)

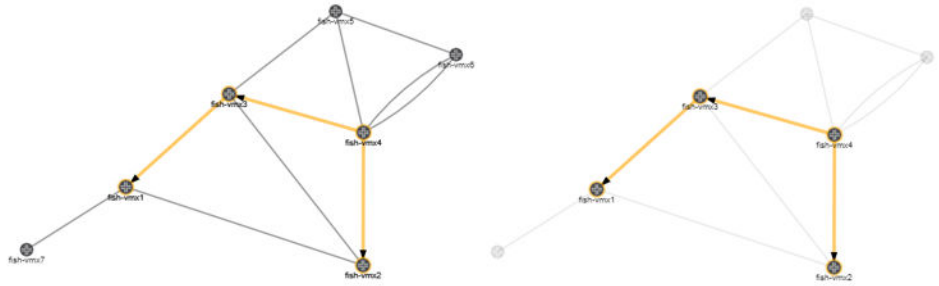
Tab	Description
	<p data-bbox="462 352 885 388">Figure 53: Opacity Effects: Example</p>  <p>The figure consists of two network diagrams side-by-side. Both diagrams show a network of nodes and edges. The nodes are labeled with IDs: fish-vms1, fish-vms2, fish-vms4, fish-vms5, fish-vms6, and fish-vms7. In the left diagram, the nodes fish-vms1, fish-vms2, fish-vms4, and fish-vms6 are highlighted with a yellow glow, and the edges connecting them are also highlighted in yellow. The other nodes and edges are shown in a light gray, semi-transparent style. In the right diagram, the nodes fish-vms1, fish-vms2, fish-vms4, and fish-vms6 are highlighted with a yellow glow, and the edges connecting them are also highlighted in yellow. The other nodes and edges are shown in a light gray, semi-transparent style.</p>

Table 105: Topology Settings Options (Continued)

Tab	Description
Map	<p>You can perform the following tasks:</p> <ul style="list-style-type: none"> Select a Light or Dark theme for the topology map. <p>Figure 54 on page 676 shows an example of the light and dark map styles.</p> <p>Figure 54: Light and Dark Map Styles: Example</p> <ul style="list-style-type: none"> Show Graticules: Displays graticules (a grid of lines that is parallel to the meridians of longitude and the parallels of latitude) and labeling of major populated places (both shown in Figure 54 on page 676). Show World Map: You can select the type of map (based on color or labeling) to be displayed on the Topology page. Based on your selection, the topology map is updated immediately. You can select one of the following options: <ul style="list-style-type: none"> PathFinder—Displays the topology on the greyscale world map. PathFinder No Labels—Displays the topology on the world map without highlighting the city or country names. Thunderforest Atlas—Displays the topology map on a colored world map with city and country names highlighted.

RELATED DOCUMENTATION

| [About the Topology Page](#) | 677

About the Topology Page

IN THIS SECTION

- [Navigation on the Topology Page | 678](#)

The Topology (**Network > Topology**) page is the main work area for the live network you load into the system. You can configure various network elements in your topology such as nodes, links, and tunnels from the network information table.

[Figure 55 on page 677](#) shows the all the components on the Topology page.

Figure 55: Topology Page

The screenshot displays the Paragon Automation interface for network topology management. Key components include:

- Left Widget:** A sidebar containing network metrics and performance data.

Metric	Value	Percentage
Active Paths	42 of 42	100%
Active Links	12 of 12	100%
Active PCC	7 of 9	78%
- Interactive Topology Map:** A central diagram showing network nodes and their connections. Nodes include ios-xr9, vmx104, vmx105, vmx106, vmx107, vmx108, and vmx109.
- Topology Menu Bar:** A vertical bar on the right side of the map containing various interaction icons.
- Network Information Table:** A table at the bottom providing detailed information for selected nodes.

Hostname	IP Address	Type	PCEP Status	NETCONF Status	AS	ISIS Area	Management IP	PCEP IP	Last Update
ios-xr9	10.0.0.109		Up		11	490011		172.16.18.109	

The Topology page consists of the following:

- **Left Widget**— Movable list of network metrics and performance data. You can hover over the More Options (horizontal ellipsis) icon to display network metrics or performance submenu. Your selections are reflected in the topology map where applicable. You can expand or collapse the left widget by clicking the arrow icon. You can drag-and-drop this widget to position it anywhere on the screen. For more information, see "[Left Widget Options on Topology Page](#)" on page 680.

- **Interactive Topology Map**—Use the topology map to access network element information and further customize the map display. You can click the legend on the bottom right corner to configure the color to be displayed for specific network elements like links. On the **Legend Settings** pane, you can click on the color to customize it and move the sliders to customize the colors for specific link utilization (%) intervals. For more information, see "[Interactive Map Features Overview](#)" on page 660.

NOTE: Topology map is refreshed whenever there are network events such as link down or up, or new node discovery. To reload the topology map, click the **Refresh** icon at the bottom left corner of the network information table.

- **Network Information Table**— Displays detailed information about network nodes, links, tunnels, interfaces, anycast groups, container LSPs, maintenance events, P2MP groups, SRLG/Facility, and topology groups within different tabs. For more information, see "[Network Information Table Overview](#)" on page 689.

NOTE: Click the collapsible arrow icon at the bottom-left of the Topology page to show or hide the network information table.

- **Topology Menu Bar**—A vertical bar at the top-right corner of the Topology page, which consists of the following:
 - Target icon—Center the topology map in the window.
 - Plus icon—Zoom in (enlarge) the topology map.
 - Minus icon—Zoom out (reduce) the topology map.
 - Settings icon—Access the topology settings window. For more information, "[Topology Settings Pane](#)" on page 669.

Navigation on the Topology Page

[Table 106 on page 678](#) describes the Topology page navigation functions.

Table 106: Topology Page Navigation Functions

Function	Method
Drag and drop	Select an element, drag to the required position on the screen, and then release.

Table 106: Topology Page Navigation Functions (Continued)





Function	Method
Select an element	Click a link or node to select it.
Select multiple elements	Do one of the following: <ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • Use the mouse scroll wheel. • Pinch to zoom using the touch pad. • Click the + or - buttons on the Topology Menu Bar.
Zoom to fit 	Click the target icon on the Topology Menu Bar to resize and center the topology map to fit the visible area of the Topology page.
Right-click to access functions	Right-click the blank part of the topology map or on a map element to access context menus.
Hover	Hover over a network element in the topology map to display the element name or ID.
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 the page or widget to fix it at any place on the screen.

Table 106: Topology Page Navigation Functions *(Continued)*

Function	Method
Resize panes	Click and drag any of the pane margins to resize the panes in a display.

RELATED DOCUMENTATION

[Network Information Table Overview | 689](#)

Left Widget Options on Topology Page

The left widget menu allows you to filter the data displayed on the topology map. The left widget displays **Metrics** when you first log in to the web user interface. Hover over the **More Options** (horizontal ellipsis) to view the other left widget options.

[Table 107 on page 681](#) describes the left widget options.

Table 107: Left Widget Options

Option	Description
Metrics	<p data-bbox="467 369 1390 432">Displays the percentage and count of the network's active paths, active links, and active PCCs that are in an UP state.</p> <p data-bbox="467 464 1417 564">The display is updated every one to two minutes, depending on the frequency of incoming events. The busier the network, the more frequent is the update. The last updated timestamp is displayed at the bottom of the widget.</p> <p data-bbox="467 596 1386 697">The number of paths detoured (using a bypass LSP) and LSPs that are being provisioned are also displayed. These numbers could differ from what is reported in the network information table.</p> <ul data-bbox="467 728 1365 758" style="list-style-type: none"> <li data-bbox="467 728 1365 758">• Active Paths: Displays the number and percentage of active paths in the topology. <p data-bbox="505 789 1409 959">By design, the Active Paths reported is not the same as what is reported in the Tunnel tab of the network information table because the Tunnel tab includes secondary paths and the Active Paths display does not. If you have a secondary path for any LSPs, the Active Paths displayed and the information in the Tunnel tab in the network information table do not match.</p> <ul data-bbox="467 995 1349 1024" style="list-style-type: none"> <li data-bbox="467 995 1349 1024">• Active Links: Displays the number and percentage of active links in the topology. <p data-bbox="505 1056 1417 1335">Active link numbers always match the Link tab in the network information table if the internal model is in sync with the live network. A mismatch indicates that the internal model has become out of sync with the live network. On a regular basis, when the internal model is updated, it is with changes to the live network topology, not with a rebuilding of the entire topology. So over time, the model and the live network can become out of sync. To correct this problem, reload the internal model with the entire live network information using Reset/Sync Network Model (Configuration >Network > PathFinder). For more information, see Pathfinder Settings on page 206.</p> <ul data-bbox="467 1371 1341 1400" style="list-style-type: none"> <li data-bbox="467 1371 1341 1400">• Active PCC: Displays the number and percentage of active PCC in the topology. <p data-bbox="505 1432 1414 1640">By design, the Active Path Computation Client (PCC) reported is not the same as what is reported in the Node tab of the network information table because the Node tab includes pseudo nodes and the Active PCC display does not. The Active PCC display only includes nodes that are routers; it does not include pseudo nodes such as Ethernet nodes or AS nodes. If you have pseudo nodes in the network, the Active PCC display and the Node tab in the network information table do not match.</p>

Table 107: Left Widget Options (Continued)

Option	Description
Performance	<p>Display current (live network) or historical (analytic traffic collection) data on the topology map.</p> <ul style="list-style-type: none"> For Current options, you can select any one of the options that you want display on the topology map. <p>NOTE: The colors are displayed based on the legend settings configured (on the bottom-right corner of the page).</p> <ul style="list-style-type: none"> For the Historical options, a slider is displayed in the upper left corner of the topology page. You can slide this to view the historical topology data based on the time range configured in the Settings.

RELATED DOCUMENTATION

[About the Topology Page](#) | 677

Group Nodes

You can represent a collection of nodes on the topology map as a single entity called node groups.

To create a node group:

1. Click **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. Select the nodes on the topology map by holding down the Shift key and left mouse button and dragging the mouse to create a rectangular selection box. All elements within the box are selected.

3. Right-click and select **Grouping > Group Selected Nodes**.

The Group Selected Nodes window appears.

4. Enter a name in the **Group Name** field.

NOTE: The group name should be alphanumeric and unique. It cannot be the same as any node ID or hostname.

5. Click **OK**.

The selected nodes are grouped and displayed as a single entity on the topology with a new Node Group icon highlighted in yellow.

The new group name is displayed in the **Grouping > Nodes & Groups** page when you right-click anywhere on the screen.

On the **Nodes & Groups** page, you can control how the group is displayed in the topology map. You can uncheck the Group Name checkbox to hide the node group on the topology map. When you expand a group in the Nodes & Groups list, the nodes are displayed individually on the map. When you collapse a group in the Nodes & Groups list, only the group name appears on the topology map represented by an icon.

NOTE: You can also double-click the Group Node icon to view the nodes separately on the topology map.

RELATED DOCUMENTATION

| [Interactive Map Features Overview](#) | 660

Group Nodes and Links into a Topology Group

You can group the nodes and links on a topology map to form a topology group. For more information, see "[About the Topology Group Tab](#)" on page 801.

To add a topology group:

1. Click **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. Select the required nodes or links on the topology map by holding down the Shift key and the left mouse button, and then dragging the mouse to create a rectangular selection box.

All elements within the box are selected.

3. Right-click and select **Grouping > Group Selected Nodes/Links into Topology Group**.

The Group Selected Nodes/Links into Topology Group window appears.

4. Enter a positive integer number in the **Slice ID** field.

The slice ID can take a value between 1 and $2^{*}53-1$. If the slice ID is assigned a value greater or equal to $2^{*}32$, the slice ID and the topology group ID might not be the same. So, we recommend that you use a value lesser than $2^{*}32$ for a slice ID.

5. Click **OK**.

The topology group is added to the **Topology Group** tab of the network information table.

RELATED DOCUMENTATION

[Interactive Map Features Overview | 660](#)

[About the Topology Group Tab | 801](#)

Ungroup Nodes

Nodes that are part of a node group are displayed as a single entity (node group icon) on the topology map. To view these nodes as discrete entities on the topology map, you can ungroup the nodes from the node group.

To ungroup nodes from an existing node group:

1. Click **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. Right-click on the node group and select **Grouping > Ungroup Selected Nodes**. Alternatively, you can double-click the node group icon (highlighted in yellow) to ungroup the nodes on the topology map.

The nodes within the node group are ungrouped and are visible as separate nodes on the topology map. The node group is removed from the Nodes & Groups page (right-click anywhere on the **Topology** map and select **Grouping > Nodes & Groups**).

RELATED DOCUMENTATION

[About the Topology Page | 677](#)

Automatically Group Nodes

Auto Grouping option enables you to use multiple rules in sequence to group nodes.

To auto group nodes:

1. Click **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. Select the nodes on the topology map by holding down the Shift key and left mouse button and dragging the mouse to create a rectangular selection box. All elements within the box are selected.
3. Right-click and select **Grouping > Auto Grouping**.
Auto grouping pane is displayed on the right.
4. Hover over the **Add** list and select a rule type from one of the following:
 - City
 - Country
 - AS
 - ISIS Area
 - OSPF Area
 - Site
 - Regular Expression: Add a more specific rule by using regular expressions to group nodes by Hostname, IP address, or Type. You can find the first match for any case-sensitive expression.

NOTE:

- You can create one rule for each rule type.
- The Edit (pencil icon) function is only available for Regular Expression rules.
- You can change the order of the rules by clicking on a rule by selecting **Up** or **Down** from the **Move** list to reposition the rule in the list.
- You can also select the check box (at the bottom of the page) to apply auto-grouping to the nodes that you have selected on the topology map. By default, this is disabled.

5. Click **Submit**.

NOTE: To delete a rule, select the rule and click the Delete (trash can) icon. On the confirmation message, click **Yes**.

RELATED DOCUMENTATION

| [Interactive Map Features Overview](#) | 660

Manage Map Layouts

From the **Map View** page, you can save topology map layouts, quickly load them to the topology map, edit the saved layouts, or delete saved layouts.

To save a map layout so you can quickly load it into the topology map:

1. Click **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. Right-click on the topology map and select **Layout > Manage Map View.**

The Manage Map View pane is displayed on the right.

3. Click **Save.**

4. Enter a name for the current layout.

5. Enter a description for the current layout.

6. Specify whether the saved layout is to be shared with all operators (shared) or is to be available only to you (private).

7. Click **Save.**

The current layout is saved and listed on the Map View table.

NOTE: You can also right-click a blank part of the topology map pane and select **Save Default Map Layout** to save the current layout as your default. This action saves the current layout as your default, but does not change the name of the default map layout in the **Manage Layouts** window.

[Table 108 on page 686](#) describes the functions you can perform on the saved layouts.

Table 108: Map View Window Buttons

Button	Functions
Load	Load the selected layout to the topology map.

Table 108: Map View Window Buttons *(Continued)*

Button	Functions
Edit	<ul style="list-style-type: none"> • Edit the name of the selected layout by clicking Edit > Rename. In the New Name field, enter the new name of the selected layout and click OK. • Edit the description of the selected layout by clicking Edit > Change Description. In the New Description field, enter the new description for the selected layout and click OK. • Set the selected layout as default by clicking Edit > Set to default. <p>NOTE: The default layout is displayed initially whenever you log in to Paragon Automation.</p>
Delete	<p>Select a layout from the list and click the delete (trash can) icon. On the confirmation message, click Yes. The layout is deleted and removed from the Map View list.</p>

NOTE: You can also show/hide columns in the table by clicking the **More Options** (vertical ellipsis) icon and selecting the specific tabs you want to display.

RELATED DOCUMENTATION

[Interactive Map Features Overview](#) | 660

Network Information Table

IN THIS CHAPTER

- [Network Information Table Overview | 689](#)
- [About the Node Tab | 692](#)
- [Add a Node | 696](#)
- [Edit Node Parameters | 699](#)
- [Delete a Node | 701](#)
- [About the Link Tab | 702](#)
- [Add a Link | 705](#)
- [Edit Link Parameters | 708](#)
- [Delete a Link | 709](#)
- [About the Tunnel Tab | 710](#)
- [Understand How Pathfinder Handles LSPs | 715](#)
- [Reroute LSPs Overview | 718](#)
- [Segment Routing Overview | 719](#)
- [Add a Single Tunnel | 729](#)
- [Add Diverse Tunnels | 743](#)
- [Add Multiple Tunnels | 754](#)
- [Edit and Delete Tunnels | 764](#)
- [About the Demand Tab | 765](#)
- [About the Interface Tab | 767](#)
- [Container LSP Overview | 768](#)
- [About the Container LSP Tab | 769](#)
- [Add a Container LSP | 770](#)
- [Edit Container LSP Parameters | 775](#)
- [Maintenance Event Overview | 775](#)
- [About the Maintenance Tab | 777](#)
- [Add a Maintenance Event | 778](#)

- [Edit a Maintenance Event | 780](#)
- [Simulate a Maintenance Event | 781](#)
- [Delete a Maintenance Event | 782](#)
- [About the P2MP Groups Tab | 783](#)
- [Add a P2MP Group | 791](#)
- [Edit P2MP Group Parameters | 798](#)
- [About the SRLG/Facility Tab | 799](#)
- [Add an SRLG/Facility | 800](#)
- [Edit SRLG/Facility Parameters | 801](#)
- [About the Topology Group Tab | 801](#)
- [Add Anycast Group Tunnels | 803](#)

Network Information Table Overview

IN THIS SECTION

- [Tasks You Can Perform | 690](#)

The Network Information Table at the bottom of the Topology page displays detailed network information based on one of the following tabs selected:

- **Node**—View node information and add, edit, or delete nodes in the network.
- **Link**—View link information and add, edit, or delete links in the network.
- **Tunnel**—View tunnel information and provision, edit, or delete tunnels in the network.
- **Interface**—View information about various interfaces (such as IPv4, MAC, and VRF) in the network.
- **Anycast Group**—View information about different anycast groups in the network. You can view the anycast group **Prefix Address**, **State** (valid or not), **SR** (segment routing) information such as flags and index, and the **Members** of the anycast group.

Segment Routing is a forwarding architecture which instructs a router on what to do with certain packets. Multiple types of SIDs supported such as node SIDs, prefix SIDs, Adjacency SIDs, Binding SIDs along with Anycast SIDs.

Anycast SIDs are a type of prefix SID that represents a group. It is present in multiple devices, and the network can reach any of the members on the group, based on IGP shortest path or potentially any other constraint defined. Anycast SID have multiple use cases, for example, they can be used as transit SIDs (loose hops) on a SR-TE policy, or used as destination (representing a service, or representing border gateways on a multi-domain network).

For LSPs using anycast group as intermediate hop, the intermediate hop will be a set of nodes. Within an anycast group, all the routers advertise the same prefix with the same SID value, which facilitates load balancing.

- Container LSP—View container LSP information and add, edit, or delete container LSPs and their sub-LSPs.
- Maintenance—View information about existing maintenance events and add, edit, simulate, or delete these events.
- P2MP Group—View information about P2MP group and their sub-LSPs, and add, edit, or delete them.
- SRLG/Facility—View information about shared risk link group (SRLG) or Facilities and add, edit, or delete them.
- Topology Group—View information about topology groups in the network. You can edit or delete a topology group.

The **Node**, **Link**, and **Tunnel** tabs are displayed by default. To view other tabs, you can hover over the **More** Tabs list (next to the **Tunnel** tab) and select required tab.

NOTE: You can show or hide the network information table by clicking the collapsible arrow icon.

Tasks You Can Perform

You can perform the following tasks:

- Related to nodes, see ["About the Node Tab" on page 692](#).
- Related to links, see ["About the Link Tab" on page 702](#).
- Related to tunnels, see ["About the Tunnel Tab" on page 710](#).
- Related to interfaces, see ["About the Interface Tab" on page 767](#).

- Related to container LSPs, see ["About the Container LSP Tab" on page 769](#).
- Related to maintenance events, see ["About the Maintenance Tab" on page 777](#).
- Related to P2MP groups, see ["About the P2MP Groups Tab" on page 783](#).
- Related to SRLG/Facility, see ["About the SRLG/Facility Tab" on page 799](#).
- Related to topology groups, see ["About the Topology Group Tab" on page 801](#).
- Filter Entries—Filter the table entries by adding new filtering criteria.

Hover over the **Filter** (funnel) icon and select **Add Filter**. On the **Add Criteria** page, select the filtering criteria from the **Field** and **Condition** list, and enter the text to be compared in the **Value** field. Click **Add**.

The filtered table entries are listed and the filter criteria name is displayed above the table column names.

You can click the cross (X) icon (next to the filter name) to remove the filtering criteria.

NOTE:

- You can add multiple filtering criteria. Once you add the multiple filtering criteria, you can select the **And** condition to display the entries matching all the filtering criteria or select the **Or** condition to display the entries matching any one of the filtering criteria.
- Quick filter: Save the filtering criteria as quick filters. Once you have added all the filtering criteria, you can save a particular criteria or multiple criteria for future use by clicking **Save**.

On the **Save Filter** window, 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 a specific column in the table.

Hover over the **More Options** (vertical ellipsis) > **Show/Hide Columns** and select the *Column-Name* check box of the columns you want to display in the table.

- Reset Preference—Reset the displayed columns to the default set of columns for each tab in the table.

Hover over the **More Options** (vertical ellipsis) and select **Reset Preference**.

- **Sort Entries**—Click the column name to highlight the up and down arrows next to the column name. Sort the table entries in ascending or descending order of that column by clicking the up or down arrow respectively.

NOTE: The network information table is not refreshed automatically. If there is a network event, a small yellow indicator next to Refresh icon at the bottom-left of the table is displayed. You must refresh the table explicitly.

RELATED DOCUMENTATION

| [About the Topology Page](#) | 677

About the Node Tab

IN THIS SECTION

- [Tasks You Can Perform](#) | 692

You can view detailed information, add, edit, or delete nodes from the **Node** tab of the network information table on the Topology (**Network > Topology**) page.

Tasks You Can Perform

You can perform the following tasks on the Node tab:

- **Add a Node.** See ["Add a Node" on page 696](#).
- **Edit Node Properties.** See ["Edit Node Parameters" on page 699](#).
- **Delete a Node.** See ["Delete a Node" on page 701](#).
- From the **View** list, you can view:
 - **Config**—View the configuration of the selected device in the network.

- **Node Traffic**—View a graphical representation of node traffic (in bps) based on the selected time range (for previous 3 hours, 1 day, 1 week, or Custom time range). A **Node Traffic** pop-up appears which you can pin anywhere on the screen.
- **Device Detail**—Displays device details (overview, alarms, and alerts), inventory details (chassis, interfaces, licenses, and software), and the configuration template. For more information, see ["View Device Statistics and Inventory information" on page 157](#).
- **Device Health**—Displays the device health information in a table and tile view on the **Network Health** page. For more information, see ["About the Network Health Page" on page 833](#).
- From the **Diagnostics** list, you can run CLI commands on the routers in the network without manually logging into the routers. You can select the routers, select the commands, specify various command parameters, execute the commands, and view/save the results. This is a unified way to manage ping and traceroute results, and is a useful tool for troubleshooting. Juniper, Cisco, Alcatel, and Huawei command sets are provided by default, and you can add other vendor command sets as needed.
- **Ping/Traceroute**—On the **Ping** or **Traceroute** window, **Default** and **Custom** tabs are displayed.

On the Default Tab:

- Click the **From** list and select one or more nodes from which ping or traceroute must be initiated.
- Click the **To** list and select the one or more nodes as destination.
- (Optional) You can click the **Use Management IP Address** check box. If you don't opt to use the management IP address, the loopback address is used.
- (Optional) You can use the **Advanced Options** to customize the ping or traceroute command. See [Table 109 on page 693](#).

NOTE: We recommend that you do not use the Advanced Options if you are running ping or traceroute for a large number of devices as these actions would be significantly slower.

Table 109: Advanced Options for Ping and Traceroute

Ping

Pattern	Enter the fill pattern (type of bits contained in the packet). You can set the bits to all ones, all zeros, a random mixture of ones and zeros, or a specific hexadecimal data pattern that can range from 0x0 - 0xFFFFFFFF. The default data pattern is all zeros.
---------	---

Count	Enter the number of ping requests to be sent.
Size	Enter the size of the ping request packets. Range: 0 through 65468 bytes.
TOS	Enter the IP type-of-service value. Range: 0 through 256.
Traceroute	
Wait	Enter the maximum wait time (in seconds) after sending the final packet.
TTL	Enter the IP maximum time-to-live value.

On the Custom tab:

- From the list at the top, you can select one of the following command categories for both Ping and Traceroute:
 - MPLS/TE Commands
 - TE/TP Commands
 - P2MP Commands
 - NIL FEC Commands
 - SR Commands

Ping offers an additional **General Commands** category.

- On the **List of Commands** window, select the **Select all** check box to select all the commands in the category. Otherwise, select the check boxes beside the individual commands of your choice.
- Once you select a command that requires a value for variable parameters, a field for each parameter is displayed under the **Selected Commands** section at the bottom of the page. Enter the appropriate values.
- Click **Submit** to execute the command(s).

The Diagnostics window displays the new commands along with status and results. When a traceroute command is successfully completed, the path is highlighted in the topology map.

- Run CLI—On the **Run CLI** window, select a command category from the list and enter the appropriate parameter values. Click **Submit** to execute the command.

NOTE: Make all your command selections first, and then enter the parameter values because the Selected Commands section of the page refreshes when you add commands and clears any parameter information already entered.

The Diagnostics window displays the new commands along with status and results. If you selected multiple nodes, each command you specified is run on each node and all the results displayed as the **Results** on the Diagnostics window.

NOTE: The selected nodes must be of the same vendor because only one CLI command set is used. If you want to run CLI on nodes of different vendors, run them separately.

- Download node information—Click **Download** to download detailed information about all the existing nodes in the topology in CSV format.
- From **More List**, you can:
 - View details about the node by clicking the Details icon when you hover over the node name or click **More > Show Detail**. A moveable pop-up with traffic (for previous 3 hours, 1 day, 1 week, or Custom time range) and node details is displayed which you can pin anywhere on the screen.
 - Filter the selected node on the Topology Map. Only the selected node is displayed.
 - Zoom in to the selected node on the Topology Map.
 - Request NETCONF reconnect—If the NETCONF status of the node is Down, you can request for a reconnect. A confirmation message appears stating that the NETCONF reconnect request is submitted successfully. If the reconnect is successful, the status is changed to Up.
 - Run Device collection for the selected node. The device collection task is added to the Task Scheduler where you can view the summary, status, and history of tasks. For more information, see ["Add a Device Collection Task" on page 978](#).
 - Force Delete—Delete one or more nodes from the topology. Force delete is used to delete nodes that are not completely withdrawn. When the TopoServer has not received the complete node withdrawal message from the network for the node, the node withdrawal is considered incomplete. This can cause the network model (maintained by TopoServer) to be out-of-sync. You can fix this sync issue by force deleting the node and then syncing the network model again.



CAUTION: Force delete a node only in extraordinary circumstances, such as troubleshooting, because force deleting a node that's working normally might cause Paragon Automation to be out-of-sync with the live network.

To force delete a node:

1. Select one or more nodes from the **Node** tab of the network information table and click **More > Force Delete**. Alternatively, right-click on the selected nodes and select **Force Delete**.

A confirmation message appears.

2. Click **Yes**.

The node is deleted from the table and the topology map. After the node is deleted, synchronize the network model. For more information, see "[Modify Pathfinder Settings From the GUI](#)" on page 206.

RELATED DOCUMENTATION

[About the Topology Page](#) | 677

[Group Nodes](#) | 682

Add a Node

You can add customer edge nodes, provider edge nodes, and nodes representing Test Agents [Paragon Active Assurance (PAA) agents] and sites (multiple nodes in one location) in the Node tab of the topology page.

To add a new node to your network:

1. Click **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. On the **Node** tab, click the add (+) icon.

The Add Node page appears on the right.

3. Configure the node parameters as per [Table 110 on page 697](#).

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

4. Click **Add**.

A confirmation message is displayed stating that the add node request is sent successfully. The new node is displayed under the **Node** tab.

Table 110: Fields on the Add Node Page

Field	Descriptions
Properties	
Name	Enter a unique name for the node.
Longitude	<p>Enter the longitude value of the node. Longitudes range from -180 to 180.</p> <p>NOTE:</p> <ul style="list-style-type: none"> • Positive longitudes are east of the Prime Meridian and negative values (precede with a minus sign) are west of the Prime Meridian. • You can either enter the values directly or use the up and down arrows to increment or decrement the values.
Latitude	<p>Enter the latitude value of the node. Latitudes range from -90 to 90.</p> <p>NOTE:</p> <ul style="list-style-type: none"> • Positive values of latitude are north of the equator and negative values (precede with a minus sign) are south of the equator. • You can either enter the values directly or use the up and down arrows to increment or decrement the values.
Site	Enter the geographical site to which the node belongs, if any. You can use specification of a geo-site to group tagged nodes in the UI.

Table 110: Fields on the Add Node Page *(Continued)*

Field	Descriptions
Role	<p>Select one of the following IP roles for the node:</p> <ul style="list-style-type: none"> • Access—For nodes/sites that are end destinations. These nodes/sites cannot be used for transit. • Core—For nodes that can be transit nodes but cannot terminate LSPs from access nodes directly. This option is reserved for future use and is not relevant to P2MP tree designs with diverse PE to CE links. • Regular—For nodes that can be used for transit. • Active Assurance—For modeling Test Agents.
Node Type	<p>Select either CE or Site. This selection depends on your network topology. In some networks, the termination is on a CE node. In others, the termination is on a network beyond the CE, which would be considered a site.</p>
Advanced	
Description	<p>Enter a description for the node.</p>
Slices	<p>Enter one or more IDs of the network slices to which you want to assign the node. Range is 1 through 2^{64} (18,446,744,073,709,551,615).</p> <p>The network slice IDs are also referred to as topology slice IDs.</p>

Table 110: Fields on the Add Node Page (Continued)

Field	Descriptions
Custom Attributes	<p>Simple name-value pairs which can be used to add any arbitrary customer-specific information. For example, to differentiate between properties for different vendor nodes.</p> <p>To add custom properties associated with the node:</p> <ol style="list-style-type: none"> 1. Click add (+) icon to add a new row. 2. Click the newly added row to enter the Name and Value. 3. Click ✓ icon to save your changes. <p>NOTE:</p> <ul style="list-style-type: none"> • You can add multiple rows (properties). • To delete an entry (row), select the row and click the delete (trash can) icon.

RELATED DOCUMENTATION

| [About the Node Tab](#) | 692

Edit Node Parameters

To edit node properties:

1. Click **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. On the **Node** tab, select the node you want to edit.
3. Click the edit (pencil) icon.

The Edit Node page appears.

4. Edit the fields as described in [Table 111 on page 700](#).
5. Click **Edit**.

A confirmation message is displayed stating that the edit node request is sent successfully.

Table 111: Fields on Edit Node page

Field	Description
Properties	
Name	Edit the node name.
Longitude	<p>Edit the longitude value of the node. Longitudes range from -180 to 180.</p> <p>NOTE:</p> <ul style="list-style-type: none"> • Positive longitudes are east of the Prime Meridian and negative values (precede with a minus sign) are west of the Prime Meridian. • You can either enter the values directly or use the up and down arrows to increment or decrement the values.
Latitude	<p>Edit the latitude value of the node. Latitudes range from -90 to 90.</p> <p>NOTE:</p> <ul style="list-style-type: none"> • Positive values of latitude are north of the equator and negative values (precede with a minus sign) are south of the equator. • You can either enter the values directly or use the up and down arrows to increment or decrement the values.
Site	Edit the geographical site name to which the node belongs, if any.
Support Secondary Path	<p>Toggle the button to enable or disable support for a secondary path.</p> <p>NOTE: This field is enabled by default.</p>
Allow any SID at First Hop	<p>Toggle the button to allow SID at first hop. When disabled, the first hop is an adjacency SID, even if the LSP is configured to use a node SID as the first hop. If enabled, the ingress node supports any SID as the first hop of the SR LSP. In this case, a node SID can be used as the first hop. This is supported on PCC devices running Junos OS Release 18.3 or later, and requires the configuration of <code>set protocol source-packet-routing inherit-label-nexthops</code>.</p>
Advanced	

Table 111: Fields on Edit Node page (*Continued*)

Field	Description
Description	Edit the description for the node.
Slices	Edit the topology slice IDs (positive integer value). Range is 1 through 18446744073709551615.
Custom Attributes	Edit the custom properties associated with the node.
Extra IP Addresses	<p>Add destination IP addresses in addition to the default IPv4 router IP address, and assign a descriptive tag to each. You can then specify a tag as the destination IP address when provisioning an LSP.</p> <ol style="list-style-type: none"> 1. Click add (+) icon to add a new row. 2. Click the newly added row to enter the Tag Name and IP address. 3. Click ✓ icon to save your changes. <p>NOTE:</p> <ul style="list-style-type: none"> • You can add multiple rows. • To delete an entry (row), select the row and click the delete (trash can) icon.

RELATED DOCUMENTATION

| [Add a Node](#) | 696

Delete a Node

You can delete a node from the topology when the node is decommissioned from the network. Since the nodes are decommissioned, deleting the nodes has no effect on the existing network. Pathfinder clears the application data associated with the deleted node (such as topology display and license count).

You can delete a node from the topology only if:

1. The node is isolated, which means all the links associated with the node are deleted. For information about deleting links, see ["Delete a Link" on page 709](#).
2. The node does not have IS-IS, OSPF, or Path Computation Element Protocol (PCEP) adjacencies. Pathfinder clears the IS-IS and OSPF adjacencies when Toposerver receives a NODE WITHDRAW message from BGP Monitoring Protocol (BMP) and and clears the PCEP adjacencies when the PCEP session is terminated.

To delete a node from the topology:

1. Click **Network > Topology**.

The Topology page appears.

2. On the **Node** tab, select the node that you want to delete.

3. Click the delete (trash can) icon.

A confirmation message appears.

4. Click **Yes**.

The node is deleted from the topology map and is removed from the network information table.

RELATED DOCUMENTATION

| [About the Node Tab | 692](#)

About the Link Tab

IN THIS SECTION

- [Tasks You Can Perform | 702](#)

You can view detailed information, add, edit, or delete links from the **Link** tab of the network information table on the Topology (**Network > Topology**) page.

Tasks You Can Perform

You can perform the following tasks on the Link tab:

- Add a Link. See ["Add a Link" on page 705](#).

- Edit Link Properties. See ["Edit Link Parameters" on page 708](#).
- Delete a Link. See ["Delete a Link" on page 709](#).
- From the **View** list, you can view:
 - Link Events—View events associated with the link based on the selected time range.

On the **Events** window, click the **Start Date to End Date** field to select the start/end date and time on the calendar for which you want to view the link events. You can also quickly view the link events that occurred **Today**, **Past 1h** (one hour), **Past 1d** (one day), and **Past Wk** (week) by selecting the respective buttons at the bottom of the calendar.

NOTE: You can also enter the date and time manually.

The time series bar chart shows the hourly aggregated Status Up or Down counts. The Up or Down event is sent by the server to the client based on the date-time range configured.

The detailed events are displayed in the table at the bottom on the page. For example, you can view the link **Action** (add, remove, update, state change, or query) and the link **Status** (up, down, planned, or unknown).

You can also export all the events data to a CSV file by clicking **Download**.

NOTE: You can also:

- Show/Hide Columns—Choose to show or hide a specific column in the table.
Hover over the **More Options** (vertical ellipsis) > **Show/Hide Columns** and select the *Column-Name* check box of the columns you want to display in the table.
- Reset Preference—Reset the displayed columns to the default set of columns in the table.
Hover over the **More Options** (vertical ellipsis) and select **Reset Preference**.

- 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.
- 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.

- From the **Diagnostics** list, you can:
 - Show Interface—Runs the **show interfaces** CLI command for each interface associated with the link. Interface related information is displayed on the **Diagnostics** window. You can view the status, type, node for which the command is run (From), description, and time of last execution, for each associated interface. Once the commands are run, you can view detailed execution details under the **Result** tab.

NOTE: On the **Diagnostics** page, from the **New** list, you can run a new ping and traceroute command. You can also export all the diagnostic related information to a text file by selecting the command and clicking **Download**.

- Download links information—Click **Download** to download detailed information about all the existing links in the topology in CSV format.
- From **More** List, you can:
 - View details about the link by clicking the Details icon when you hover over the link name or click **More > Show Detail**. A moveable pop-up with interface stats for Node A and Node Z, traffic (for previous 3 hours, 1 day, 1 week, or Custom time range, and link details is displayed which you can pin anywhere on the screen.
 - Filter selected link on the Topology Map. Only the selected link is displayed.
 - Zoom in to the selected link on the Topology Map.
 - Trigger LSP Optimization for the selected link. For more information, see .
 - Force Delete—Delete one or more links from the topology. Force delete is used to delete links that are not completely withdrawn. When the TopoServer has not received the link withdrawal message from the network, the link withdrawal is considered incomplete. This can cause the network model (maintained by TopoServer) to be out-of-sync. You can fix this sync issue by force deleting the link and then syncing the network model again.



CAUTION: Force delete a link only in extraordinary circumstances, such as troubleshooting, because force deleting a link that's working normally might cause Paragon Automation to be out-of-sync with the live network.

To force delete a link:

1. Select one or more links from the **Link** tab of the network information table and click **More > Force Delete**. Alternatively, right-click on the selected links and select **Force Delete**.

A confirmation message appears.

2. Click **Yes**.

The link is deleted from the table and the topology map. After the link is deleted, synchronize the network model. For more information, see "[Modify Pathfinder Settings From the GUI](#)" on [page 206](#).

RELATED DOCUMENTATION

No Link Title

Add a Link

To add a link in the topology map:

1. Click **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. On the **Link** tab, click the add (+) icon.

The Add Link page appears.

3. Configure the link fields as per the [Table 112 on page 705](#).

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

4. Click **Add**.

A confirmation message appears stating that the request is successfully submitted. The new link is displayed under the **Link** tab.

Table 112: Fields on the Add Link Page

Field	Description
-------	-------------

Properties

Table 112: Fields on the Add Link Page (*Continued*)

Field	Description
Name	Enter a name for the link.
Node A	Click the Node A list and select the name of the ingress (source) node from the list.
Node Z	Click the Node Z list and select the name of the egress (destination) node from the list.
IP A	Enter the IPv4 address of Node A.
IP Z	Enter the IPv4 address of Node Z.
IPv6 A	Enter the IPv6 address of Node A.
IPv6 Z	Enter the IPv6 address of Node Z.
Bandwidth AZ and ZA	<p>Enter the bandwidth (in bps) for each direction (Node A to Z and from Z to A). You must enter a number immediately followed by K (Kbps), M (Mbps), or G (Gbps). For example, 10M signifying 10 Megabits per second.</p> <p>NOTE: If you enter a value without units, bps is applied.</p>
ISIS Off/On	<p>Toggle the button to enable or disable IS-IS protocol.</p> <ul style="list-style-type: none"> Enter IS-IS area level (L1, L2, or L1L2) metric. For more information about IS-IS area, see Understanding IS-IS Areas to Divide an Autonomous System into Smaller Groups. Enter the IS-IS TE metric. Range: 1 through 16,777,215. Default: Value of the IGP metric.
RSVP Off/On	Toggle the button to enable or disable the RSVP protocol. Enter the RSVP bandwidth for each direction AZ (Node A to Z) and ZA (from Z to A).

Constraints

Table 112: Fields on the Add Link Page (Continued)

Field	Description
Admin Group AZ and ZA	Select the bit-level link coloring from the list for A to Z and Z to A links.
Delay AZ and ZA	<p>Enter the delay metric for each direction in milliseconds for A to Z and Z to A links.</p> <p>This field is used only if the routing method is set to Delay Metric. When the program performs path placement and is trying to find the best route for a call, delay metrics are examined to determine the desirability/undesirability of a link. Two delay metrics are supported, one for each direction of the trunk. If the hardware does not support asymmetric delay metrics, the second delay is marked as '-'. If a delay metric is not defined for a trunk, a default delay is calculated based on propagation delay and serialization delay.</p>
Admin Weight AZ and ZA	Enter the admin weight (number) for links between Node A and Z. Range: $1-2^{31}$ (1 through 2147483648).
Util Reroute Threshold AZ and ZA	Enter the link utilization threshold (in percentage) for A to Z and Z to A links. When the link utilization exceeds the specified threshold, the controller reroutes the LSPs to reduce the link utilization within the threshold value.
Packet Loss Threshold	Enter the packet loss threshold in percentage for A to Z and Z to A links. When packet loss on a link exceeds this threshold, the link is considered unstable and rerouting of traffic to avoid the link is triggered. To achieve this, a maintenance event is created for each link, temporarily making the link unavailable for traffic.
Advanced	
Description	Enter a text description for the link.
Type	Click the Type field and select the type of link you want to create from the list.
Slices	Enter one or more numbers (positive integers) which represents the topology slice IDs. Range is 1 through 18446744073709551615.

Table 112: Fields on the Add Link Page (*Continued*)

Field	Description
Protected Link	<p>Toggle this button to enable protected link status for this link. By default, link protection is disabled.</p> <p>PCE supports preferred protected links routing constraint for packet LSPs. When this constraint is selected, PCE computes the path that maximizes the number of protected links, and therefore offers the best overall protection.</p>
Custom Attributes Node A and Node Z	<p>Simple name-value pairs which can be used to add any arbitrary customer-specific information. For example, to differentiate between properties for different vendor nodes.</p> <p>To add custom properties associated with the node:</p> <ol style="list-style-type: none"> 1. Click add (+) icon to add a new row. 2. Click the newly added row to enter the Name and Value. 3. Click ✓ icon to save your changes. <p>NOTE:</p> <ul style="list-style-type: none"> • You can add multiple rows (properties). • To delete an entry (row), select the row and click the delete (trash can) icon.

RELATED DOCUMENTATION

| [About the Link Tab | 702](#)

Edit Link Parameters

To edit an existing link in the topology:

1. Click **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. On the **Link** tab, select the link that you want to edit.

3. Click the edit (pencil) icon.
The Edit Link page appears.
4. Edit the fields as per [Fields on the Add Link Page on page 705](#).

NOTE: You cannot edit the Name, Node A, and Node Z fields.

5. Click **Edit**.
A confirmation message is displayed stating that the request was submitted successfully.

RELATED DOCUMENTATION

| [Add a Link | 705](#)

Delete a Link

You can delete a link from the topology when the link is decommissioned from the network. Since the links are decommissioned, deleting the links has no effect on the existing network. Pathfinder clears the application data associated with the deleted link (such as topology display and license count).

You can delete a link only if:

1. The link's operational status is **Down**. Pathfinder changes the operational status to **Down** when Toposerver receives the first LINK WITHDRAW message from BGP Monitoring Protocol (BMP).
2. The link does not have active IS-IS or OSPF adjacencies. Pathfinder drops IS-IS and OSPF adjacencies when Toposerver receives the second LINK WITHDRAW message from BMP.

To delete a link from the topology:

1. Click **Network > Topology**.
The Topology page appears.
2. On the **Link** tab, select the link that you want to delete.
3. Click the delete (trash can) icon.
A confirmation message appears.
4. Click **Yes**.
The link is deleted from the topology map and is removed from the network information table.

RELATED DOCUMENTATION

| [About the Link Tab | 702](#)

About the Tunnel Tab

IN THIS SECTION

- [Tasks You Can Perform | 710](#)

Use the Tunnel tab to view and manage tunnels (label-switched paths or LSPs).

To access this tab, select **Network > Topology**. The Topology page appears, with the topology map at the center and the network information table at the bottom of the page. The table displays various tabs, including the Tunnel tab.

Tasks You Can Perform

- Hide unrelated nodes and links—Select one or more tunnels and enable the **Hide unrelated** toggle button. Only the nodes and links that the selected tunnels pass through are displayed on the topology map.
- From the Provisioning list, you can perform the following tasks:
 - Add a tunnel—See ["Add a Single Tunnel" on page 729](#).
 - Add diverse tunnels—See ["Add Diverse Tunnels" on page 743](#).
 - Add multiple tunnels—See ["Add Multiple Tunnels" on page 754](#).
 - Manually reprovision tunnels—Select one or more tunnels for which provisioning has failed or the path isn't the expected path, and click **Reprovision**. The tunnels are reprovisioned and a confirmation message appears on the top of the page.

NOTE: You can reprovision only PCE-initiated and PCC-delegated tunnels.

- Set current path as explicit path—When creating the tunnel, if you've configured the routing path type for one or more secondary or standby tunnels as **Dynamic** or **Preferred** and now want to

configure this path as a strict explicit path, select the tunnel, and click **Set current path as explicit path**. A confirmation message appears on the top of the page and the routing path type is set to **Required**. When you set the routing path type as Required, the PCE considers this path as explicit. If the required path is not viable and available, the tunnel is down and the PCE doesn't compute an alternate path.

- Optimize tunnels—Select one or more tunnels for which you want to optimize the path, and click **Trigger Tunnel optimization**. The paths are optimized and a confirmation message appears on the top of the page.

NOTE: You can optimize only PCE-initiated and PCC-delegated tunnels.

- From the Delegation list, you can perform the following tasks:
 - 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**. See ["Add and Remove LSP Delegation" on page 816](#)
 - Return delegation of LSPs to the PCC—Select one or more LSPs (Control Type: **Delegated**) that were previously delegated to the PCE and select **Delegation > Return Delegation to PCC**. Control of the selected LSPs is temporarily returned to the PCC for a period of time based on the router's timer statement, and a confirmation message appears on the top of the page. The Control Type for the selected LSPs changes to **Device Controlled**.
- From the View list, you can perform the following tasks:
 - 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 **Events**. In the Events page that appears, choose the start date and end date from the calendar that is displayed. Click the **Select Time** link at the bottom-right corner of the calendar to select the time for which you want to retrieve the events. Alternatively, you can select one of the shortcuts at the bottom-left corner of the calendar to view the events for the current day, past hour, past day, or past week.

Then, click **OK**. The events for the tunnel are displayed in the Events table for the selected time range.

A graph with timeline view, which indicates the bandwidth spikes for a time period, is also displayed. You can drag the slider on the graph to select a custom time range. The events for the selected time range are displayed in the Events table.

To download the displayed data to your local system as a comma-separated values (CSV) file, click **Download**.

NOTE: The events displayed on the Events page are restricted to external communication to and from the Path Computation Element (PCE). Most of the communications internal to the PCE are captured only in the log files.

- View 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 times.
- View tunnel delay—At any given time, the 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 as the simple sum of all link delays in the tunnel path. To view tunnel delay in graphical form, select the tunnel and click **Tunnel 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 times.

NOTE: To view tunnel delay, you must first:

- Set the interval (in seconds) for the PCS to calculate tunnel latency.

You can specify the interval in the **lsp-latency-interval** field in the **Path Computation Server** section of the Pathfinder page (**Configuration > Network > Pathfinder**) on the GUI or by using the CLI (`set northstar path-computation-server lsp-latency-interval`).

- Configure routers to send JTI telemetry data and real-time performance monitoring (RPM) statistics to measure link delay. See "[Configure Routers to Send JTI Telemetry Data and RPM Statistics to the Data Collector](#)" on page 858 for details.

- From the Diagnostics list, you can perform the following tasks:
 - Run CLI commands—Select **Show Tunnels**, **Run MPLS Ping**, or **Run MPLS Traceroute**. The Diagnostics page appears, displaying the list of CLI commands that are running or that have completed. The Status column indicates whether running the command was successful or not. Select one or more rows in the table to display the results in the Results tab (in the lower part of the page).

From the Diagnostics page, you can:

- Initiate a new ping or traceroute command—From the **New** list, select either **Ping** or **Traceroute**.

Based on what you select, the Ping or Traceroute page appears. The Default tab and Custom tab are similar for Ping and Traceroute.

On the Default tab, you can:

- Select the ingress and egress nodes from the From list and To list.
- Choose to enable the Use Management IP address toggle button if the nodes have management IP addresses specified for out-of-band use.
- Choose to enable the Advance Options toggle button to configure advanced parameters. For ping request packets, you can configure hexadecimal pattern, count, size, and IPv4 type-of-service (tos). For traceroute, you can configure the maximum time (in seconds) to wait for a response to the traceroute request after sending the final packet, and the maximum time-to-live value.

On the Custom tab, you can:

- Select the category of commands from the list at the top of the tab. The commands specific to the selected category appear in the List of Commands section.

The options for both Ping and Traceroute include:

- General Commands (only for Ping)
- RSVP LSP Commands
- P2MP Commands
- SR LSP Commands
- From the list of commands displayed, select the **Select All** check box to select all the commands in the category. Otherwise, select the check boxes beside the command variations of your choice.
- If you select one or more commands from the list, the Selected Commands section appears. For commands that require the specification of variable parameters, a space for each parameter is displayed. Enter the appropriate value.
- Click **Submit** to execute the commands. The Diagnostics page displays the new commands along with the status and results. When a traceroute command is successfully completed, the path is highlighted in the topology map.
- Download details of all the tunnels—To view detailed information on all the tunnels displayed in the network information table, click **Download**. You can choose to open the comma-separated values (CSV) file with Excel or other applications, or save the file to your local system.
- From the More list, you can perform the following tasks:

- View details of the tunnel—Select a tunnel and click **Show Detail**. The Tunnel <Tunnel-Name> page appears, displaying the details (on the Details tab) and traffic on the tunnel for the time period that you selected (on the Traffic tab).
- 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 create a secondary or standby tunnel for the selected 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: This option is available only for PCE-initiated and PCC-controlled tunnels.

- Run device collection to obtain the latest information—Select a tunnel and click **Run Device Collection**. A confirmation message appears indicating that a device collection task has been added. Navigate to the Task Scheduler page (**Administration > Task Scheduler**) to view details of the task.
- Forcefully delete a tunnel—Sometimes, it is necessary to remove tunnels from the topology when deletion requests have been rejected by the devices or when a deletion request cannot be sent to the device because the device is decommissioned. Also, the internal model may be out-of-sync with the live network, so the tunnel may have been deleted from the router but may still appear in the GUI. In such cases, the Delete icon on the top-right corner of the network information table isn't available. Instead, you can select one or more tunnels that you want to delete, and click **Force Delete**.

An alert message appears, asking you to confirm the delete operation. Click **Yes**.

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

- Reload the network information table—To re-download the data model to your Web browser, click **Reload**. A REST API query is sent to the Paragon Automation server and the network information table is updated.
- Edit (modify) parameters configured for tunnels—See ["Edit and Delete Tunnels" on page 764](#).
- Delete tunnels—See ["Edit and Delete Tunnels" on page 764](#).

RELATED DOCUMENTATION

[About the Topology Page | 677](#)

[Network Information Table Overview | 689](#)

Understand How Pathfinder Handles LSPs

IN THIS SECTION

- [LSP Control Types | 715](#)
- [LSP Path Types | 716](#)
- [Protocols to Provision and Manage LSPs | 716](#)
- [LSP Routing Methods | 717](#)
- [Routing Path Types | 718](#)
- [Deletion of LSPs | 718](#)

In Paragon Pathfinder, the Path Computation Element (PCE) computes paths in the network and applies computational constraints. The PCE uses the Path Computation Element Protocol (PCEP) or NETCONF to learn about label-switched paths (LSPs) in the discovered network topology. You can view all LSPs and their attributes in the network information table on the Topology page of the Paragon Automation GUI.

In the rest of this topic, we explain the different parameters that define how Pathfinder handles LSPs.

LSP Control Types

The LSP's control type determines whether the Path Computation Client (PCC, which is the router) or the PCE maintains the operational and configuration states of the LSP.

The PCE supports the following control types for LSPs:

- **PCC-controlled LSPs (also known as device-controlled LSPs)**—You can configure such LSPs on the router either from the GUI (by selecting NETCONF as the provisioning method) or from the CLI. The LSPs are managed by the router, which maintains both the operational state and the configuration state of the LSPs. The LSPs are part of the router's configuration file.
- **PCC-delegated LSPs**—You cannot configure such LSPs directly from the GUI or from the CLI. You must first configure a PCC-controlled LSP from the GUI or from the CLI, and then delegate it to the PCE for management from the Configure LSP Delegation page (**Network > Tunnels > Configure LSP Delegation**) on the GUI or from the CLI. The router maintains both the operational state and the configuration state of the LSPs, and the LSPs are part of the router's configuration file.

If Pathfinder is down or if the LSP is not working as expected, you can delegate the LSP back to the PCC, in which case, the LSP is reclassified as PCC-controlled.

- **PCE-initiated LSPs**—You can configure such LSPs either from the GUI (by selecting PCEP as the provisioning method) or from the CLI. The PCE manages the LSPs and only the operational state is maintained in the router. The LSPs are not part of the router's configuration file.

NOTE: Although Pathfinder typically creates PCE-initiated LSPs, in the following cases, the PCE discovers such LSPs from the router:

- When a PCE-initiated LSP is created by an external controller other than Pathfinder, the PCE then discovers the LSP from the router.
- When you reset the topology in the GUI, the PCE re-learns the LSPs from the router.

LSP Path Types

Pathfinder 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.

Protocols to Provision and Manage LSPs

Pathfinder supports two protocols for provisioning and managing LSPs: PCEP and NETCONF. When you provision an LSP by using PCEP, the LSP is added as a PCE-initiated LSP. When you provision an LSP by using NETCONF, the LSP is added as a PCC-controlled LSP.

[Table 113 on page 717](#) lists the provisioning and managing actions available for each LSP control type.

Table 113: Protocols Used to Provision, Modify, or Delete LSPs

LSP Control Type	Provision LSPs	Modify LSPs	Delete LSPs
PCC-controlled LSP	NETCONF	NETCONF	NETCONF
PCC-delegated LSP	Not applicable (because you cannot directly create a PCC-delegated LSP)	PCEP	NETCONF
PCE-initiated LSP	PCEP	PCEP	PCEP

Irrespective of whether the LSPs are provisioned by using PCEP or NETCONF, Pathfinder can learn about LSPs by using PCEP or by device collection. Both PCEP and device collection discover the same LSP attributes.

For LSPs provisioned by using PCEP, reprovisioning (in case of a provisioning failure), deletion, rerouting, and path optimization for the LSPs are triggered automatically. For LSPs provisioned by using NETCONF (which are PCC-controlled LSPs and PCC-delegated LSPs), you must initiate the following tasks manually:

- If the provisioning of PCC-controlled LSPs fails (for example, when there is a commit failure or if the NETCONF session is down), you must resubmit the provisioning order manually.
- If the deletion of PCC-delegated LSPs or PCC-controlled LSPs fails, you must resubmit the deletion order manually.
- When the PCE receives an LSP down event from the network (for example, when maintenance events are scheduled for the LSP) for PCC-controlled LSPs, the PCE does not automatically recompute and reprovision a new path for the LSPs. You must manually modify the LSP to change the routing path.
- When you run path optimization, the PCE doesn't optimize PCC-controlled LSPs automatically. You must manually modify the LSP to choose an optimal routing path.

LSP Routing Methods

When you create an LSP, you can choose one of the following routing methods to specify whether the PCE should compute and provision the path or not:

- routeByDevice routing method—The LSP is provisioned with no explicit path, so the router computes the path.

- Other routing methods (default, delay, adminWeight, constant, distance, IS-IS, OSPF)—For routing methods other than routeByDevice, the PCE computes and provisions the path.

Routing Path Types

When you add an LSP, you can choose one of the following routing path types to specify how the PCE should compute the path:

- Dynamic—The PCE computes the path without imposing any path restrictions.
- Required—The PCE uses the path that you specified. If the specified path is not viable and available, the LSP status changes to Down and the PCE does not perform the computation to look for an alternate path.
- Preferred—The PCE uses the path that you specified, as long as it is viable and available. If not, the PCE computes an alternate path.

Deletion of LSPs

When an LSP is deleted from the router (and, thereby, from the network), it is automatically deleted from Pathfinder, unless it was previously modified by a user (either from the GUI or by using REST APIs). Any LSP that is modified by a user has a Persist state associated with it. Therefore, LSPs with Persist state must be deleted manually from the GUI or by using REST APIs. See ["Edit and Delete Tunnels" on page 764](#).

RELATED DOCUMENTATION

[Add a Single Tunnel | 729](#)

[Add Diverse Tunnels | 743](#)

[Add Multiple Tunnels | 754](#)

Reroute LSPs Overview

Paragon Automation uses Path Computation Element Protocol (PCEP) or Network Configuration Protocol (NETCONF) to learn about Label-switched Paths (LSPs) in the discovered network topology. You can configure Paragon Automation to automatically reroute LSPs.

The parameters that trigger LSP rerouting are *link-utilization-threshold*, *packet-loss-threshold* and *reroute-minimum-interval* that can be configured on the **Analytics** tab of the **PathFinder Settings** page

(**Configuration > Network Settings**) and *Maximum delay* in the **Constraints** tab when you add a single LSP. These parameters are applied to all links in the network.

For LSP rerouting based on link utilization (bandwidth), you can specify a minimum reroute interval (in minutes) and a link utilization threshold (%). The reroute interval is used to pace successive rerouting events. For more information, see *Analytics* in "[Modify PathFinder Settings From GUI](#)" on page 222 or *Maximum delay* in "[Add a Single Tunnel](#)" on page 737.

LSPs are rerouted when both of the following conditions are true:

- A link utilization threshold has been crossed. To avoid unnecessary network churn, Paragon Automation only considers rerouting an LSP with traffic or a bandwidth reservation when the link utilization threshold has been crossed.
- No previous rerouting processes have occurred within the defined reroute interval.

When a threshold has been crossed, LSPs with a lower priority setting and higher traffic are the first to be rerouted, before LSPs with a higher priority setting and lower traffic. If LSP traffic data is available, Paragon Automation uses it over bandwidth reservation for determining whether an LSP should be rerouted. If LSP traffic data is not available, Paragon Automation considers LSP bandwidth reservation to make the determination.

You can override global thresholds with link-specific thresholds. To set link-specific thresholds, you can edit the link parameters (constraints) for each link listed under the **Link** tab of the network information table in the **Topology** page (**Network > Topology**). For more information, see "[Edit Link Parameters](#)" on page 708.

RELATED DOCUMENTATION

| [Understand LSP Delegation and Undelegation](#) | 815

Segment Routing Overview

IN THIS SECTION

- [Segment Identifiers \(SIDs\)](#) | 720
- [Binding SIDs](#) | 723
- [Maximum SID Depth \(MSD\)](#) | 725
- [Rerouting and Reprovisioning](#) | 726

- [On-Demand Next-Hop, Intra-Domain \(Experimental Feature\) | 727](#)
- [View the Segment Routing Path | 728](#)
- [Points to Remember | 729](#)

Paragon Pathfinder supports a source-based routing technique that is known as Source Packet Routing in Networking (SPRING), also known as segment routing. In segment routing, an ingress router steers a packet through explicit paths in the network and doesn't rely on the transit nodes in the network to determine the path.

You can configure these explicit paths by adding a segment routing tunnel. See *About the Tunnel Tab* for information about how to add single, diverse, and multiple segment routing tunnels.

You can configure SPRING features on Juniper devices which run Junos OS Release 17.2R1 or later, and other vendor devices that support segment routing. Paragon Pathfinder supports both IS-IS and OSPF as the interior gateway protocols (IGPs) in the network for segment routing.

See the [Junos OS documentation](#) for details on segment routing concepts and support on Juniper devices which run Junos OS.

The following sections explain segment routing behavior in Paragon Pathfinder.

Segment Identifiers (SIDs)

Paragon Pathfinder compiles a set of paths that satisfy a policy into SIDs. These paths are prepended to a packet as instructions. The SIDs denote paths to nodes in a network. The nodes execute the instructions that are specified in the SID lists.

You can define the hops in a policy for a segment routing tunnel path in the Path tab that is displayed when you add or modify the tunnel. To define the hops, specify the sequence of nodes, links and groups of nodes that you want the path to traverse. The paths that satisfy your path policy will also satisfy your other requirements such as delay, bandwidth, and so on. Therefore, the result of segment routing path computation is a list of SIDs that instruct the network to implement paths that satisfy your policy. However, the SID list might not align exactly with the hops that are specified in the path policy.

Currently, Paragon Pathfinder supports node SIDs (associated with nodes), adjacency SIDs (associated with links), binding SIDs (associated with tunnels), and anycast SIDs (associated with anycast groups).

NOTE: Currently, Paragon Pathfinder does not support binding SIDs and anycast SIDs for SRv6 tunnels.

If both a node SID and an anycast SID are available for path computation, the Path Computation Server implements the paths that satisfy your policy as follows to route traffic according to the path policy:

- The Path Computation Server uses the anycast SID instead of the node SID if PCEP is used as the provisioning method.
- The Path Computation Server uses the node SID instead of the anycast SID if NETCONF is used as the provisioning method.

To view the SIDs in the GUI, do the following:


- Adjacency SID labels:
 - To view segment routing-MPLS adjacency SID labels in the topology map, right-click anywhere on the blank space in the map. From the list that appears, select **Link Label** and then, select **SID A::Z**. To view SRv6 SID labels for links on the map, select **Link Label** and then, select **SRv6 SID Function A::Z**.
 - In the network information table, the Link tab displays the segment routing-MPLS SID labels for links in the **SID A** and **SID Z** columns. The SRv6 SID labels are displayed in the **SRv6 SID A** and **SRv6 SID Z** columns.

[Figure 56 on page 722](#) shows an example topology map with segment routing-MPLS SID labels for links.

To view the attributes of each SID for a specific link, select a link in the Link tab. Click the **Details** icon that appears when you hover over the link or select **More > Show Detail**. In the Link - *Link Name* page that appears, navigate to **Details tab > endA (or endZ) > Protocols > SR > SIDs [n]**.

NOTE: Currently, Paragon Pathfinder supports only one segment routing-MPLS SID and one SRv6 SID per interface.

Figure 56: Adjacency SID Labels



- Node SID labels—In the network information table, the Node tab displays the segment routing-MPLS SID labels for nodes in the **SID** column and the SRv6 SID labels in the **SRv6 SID** column. Also, you can view the segment routing-MPLS SID labels in the topology map. The SRv6 SID labels are, however, not displayed on the topology map.

NOTE: You can view the SID labels only when the segment routing global block (SRGB) value is in the same range for all the nodes in the network.

To view the node SIDs of other nodes from the perspective of a particular node, select the node in the network information table and click **View > Node SIDs from Selected Node**. The node SIDs of the nodes are displayed in the topology map from the perspective of the selected node. For example, [Figure 2 on page 723](#) shows the node SIDs of the nodes as viewed from the perspective of node ios-xr8, while [Figure 3 on page 723](#) shows the node SIDs as viewed from the perspective of node vmx101. A node can have different node SID values based on the perspective of a particular node. For example, from the perspective of node ios-xr8, the node SID for node vmx102 is 80002, whereas, from the perspective of node vmx101, the node SID for node vmx102 is 1002, as highlighted in [Figure 2 on page 723](#) and [Figure 3 on page 723](#).

Figure 57: Node SIDs as Viewed from ios-xr8

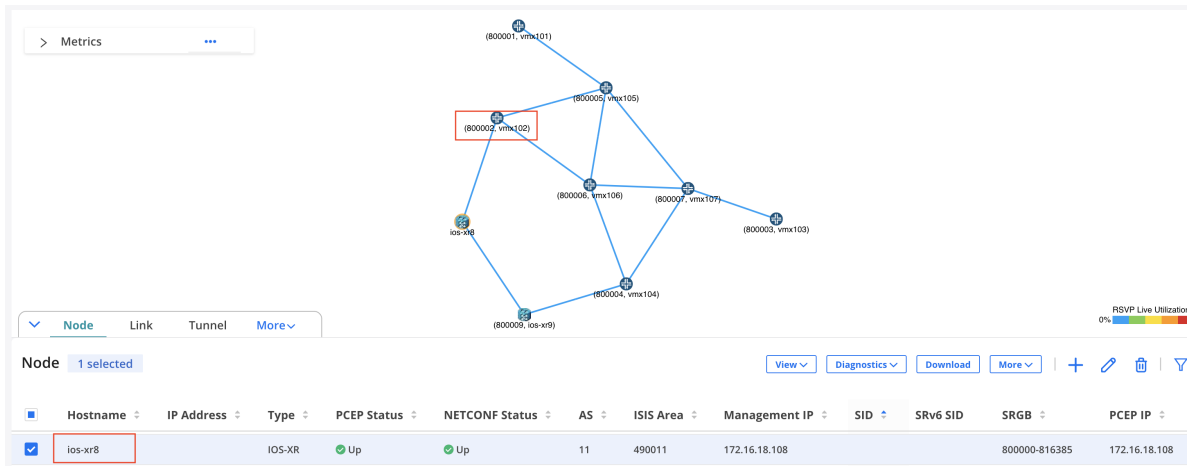



Figure 58: Node SIDs as Viewed from vmx101



- Anycast SID labels—In the network information table, the Anycast Group tab displays the anycast SID labels for segment routing-MPLS tunnels as the **Index** value in the **SR** column.
- Binding SID labels—In the network information table, the Tunnel tab displays the SID labels for binding segment routing-MPLS tunnels in the **BSID** column. The Link tab displays the SID labels in the **SID A** and **SID Z** columns.

Binding SIDs

You can provision a pair of binding SID segment routing-MPLS tunnels (one going from A to Z and one for the return path from Z to A) with NETCONF as the provisioning method. See *Add a Single Tunnel* for

details on adding a tunnel. When the tunnels are provisioned, a private forwarding adjacency is automatically created. The names of these adjacencies are in a specific format, with three sections, separated by colons. For example, **binding:0110.0000.0105:privatefa57**.

- The names all start with “binding” followed by a colon.
- The center section is the name of the originating node, followed by a colon (**0110.0000.0105:** in this example).
- The last section is the name you specified for the binding SID segment routing tunnel in the Name field in the Properties tab of the Add Tunnel page (**privatefa57** in this example).

You can tunnel only a non-binding SID segment routing tunnel over a binding SID segment routing tunnel (and vice versa). The binding reduces the number of labels in the label stack (private forwarding adjacency labels can represent multiple hops in the path).


NOTE: Currently, Paragon Pathfinder does not support binding SID label allocation or collision detection. Junos OS has built-in collision detection. Thus, if the binding SID label specified is outside the allowed range of 1000000 through 1048575, Junos OS does not allow the configuration to commit. Correspondingly, the Controller Status in the Tunnel tab of the network information table displays **FAILED(NS_ERR_INVALID_CONFIG)**.

To add a pair of binding SID segment routing tunnels, provision a second binding SID segment routing tunnel in the direction opposite to the first tunnel. You must use the same tunnel name as the first tunnel in the pair to ensure that the tunnels can be properly matched. The corresponding private forwarding adjacency link is automatically created when the tunnel pair is provisioned. The binding SID label value can also be the same as in the first tunnel in the pair, but it is not mandatory.

The private forwarding adjacency link can then be selected as a destination when you designate hops for a non-binding SID segment routing tunnel. Use show commands on the router to confirm that the tunnel pair is pushed to the router configuration.

When you select a binding SID segment routing tunnel in the network information table, the corresponding private forwarding adjacency links are displayed as dotted lines in the topology map as shown in [Figure 59 on page 725](#).

Figure 59: Private Forwarding Adjacency Links



Maximum SID Depth (MSD)

When the controller computes the segment routing paths, it must learn the MSD that it can impose at each node that corresponds to a segment routing path such that the SID stack depth value of the computed path does not exceed the number of SIDs imposed by the node.

To avoid an equipment limitation on the MSD, you can select `RouteByDevice` as the routing method when you add the segment routing tunnel with node SIDs. This option enables the router to control a part of the routing, so fewer labels need to be explicitly specified.

When you add a segment routing tunnel, a symptom of encountering the MSD limitation when you are not using `routeByDevice` is that although a row for the new tunnel is added in the network information table, the Op Status is displayed as **Unknown** and the Controller Status is displayed as **Reschedule in x Minutes**. No tunnel is created to forward the traffic. Thus, the traffic is forwarded as per the shortest path in the routing table for non-engineered traffic. To resolve this issue, you must request parameters (such as different hops) for the tunnel so that Paragon Pathfinder computes a path that does not violate the MSD of routers in the path. Alternatively, configure some tunnels with binding SIDs to create forwarding adjacencies so that Paragon Pathfinder can specify a binding SID in the SID list.

The hop information that you specify in the Path tab when you add a segment routing tunnel influences the routing. You can select hops up to the MSD hop limitation that is imposed on the ingress router, and specify **Strict** or **Loose** adherence. If you specify the hop as strict, the tunnel must take a direct path (with only the links and nodes that you specify for the 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.

Rerouting and Reprovisioning

For PCEP-provisioned segment routing tunnels, the router is able to report the operational status of only the first hop. The **Op Status** column in the network information table displays the operational status. After the first hop, the controller takes responsibility for monitoring the SID labels, and for reporting the operational status. If the labels change or disappear from the network, the controller tries to reroute and re-provision the tunnels that are in a non-operational state.

If the controller cannot find an alternative routing path that complies with the constraints, the tunnel is deleted from the network. However, these tunnels are not deleted from the data model (the tunnels persist in the data storage mechanism). The goal is to minimize traffic loss from non-viable segment routing tunnels by deleting the tunnels from the network. When a segment routing tunnel is deleted, the Op Status column displays the status as **Unknown**. The Controller Status column displays the status as **No path found** or **Reschedule in x Minutes**.

You can mitigate the risk of traffic loss by creating a secondary path for the tunnel with fewer or more relaxed constraints. If the tunnel does not meet the original constraints, the controller first tries to reroute using the secondary path. If the rerouting works, the tunnel remains **Up** and is not deleted.

NOTE: The controller permits adding a secondary path to a segment routing tunnel. However, it is not provisioned as a secondary path to the PCC because the segment routing tunnel protocol does not support secondary paths.

If you are creating a segment routing tunnel by using REST APIs, you can set the rerouting behavior of the tunnel to **noRerouting**. The segment routing Path Computation Server brings down the segment routing tunnel if topology changes cause traffic in the tunnel to deviate from the path it was originally provisioned on. If **noRerouting** is not specified (that is rerouting is allowed), the Path Computation Server computes a new path that complies with the user-defined path policy when the network topology changes. When the Path Computation Server has an option to use both node SIDs and anycast SIDs for implementing a path policy, the Path Computation Server uses SIDs as follows:

- If **noRerouting** is specified, node SIDs are used for implementing the path policy.
- If rerouting is allowed, anycast SIDs are used for implementing the path policy. When the topology changes, the existing SID list can still implement a path policy by shifting the tunnel from one node in the anycast group to another node in the anycast group.

On-Demand Next-Hop, Intra-Domain (Experimental Feature)

NOTE: This feature is for lab and demonstration purposes only. We do not recommend using this feature for production networks.

The On-Demand Next-hop (ODN) feature enables the controller to dynamically create segment routing-traffic engineering (SR-TE) tunnels when routes resolve over BGP next hops. The SR-TE tunnels can then be delegated and managed by the Path Computation Element (PCE). To use this feature, configure the device with a recent version of Junos OS 20.4 or later that supports segment routing ODN (for example, Junos 20.4I-20200910).

You must configure the Junos OS devices to create the tunnels. The following example shows the prerequisite configuration:

```
set routing-options dynamic-tunnels odncf spring-te source-routing-path-template
odnmytemplate
set routing-options dynamic-tunnels odncf spring-te destination-networks 10.0.0.11/32
set protocols source-packet-routing compute-profile test-compute-prof
no-label-stack-compression
set protocols source-packet-routing compute-profile test-compute-prof
maximum-computed-segment-lists 1
set protocols source-packet-routing source-routing-path-template odnmytemplate
lsp-external-controller pccd
set protocols source-packet-routing source-routing-path-template odnmytemplate primary
test-computer compute test-compute-prof
```

In the *routing-options dynamic-tunnels configuration* section, specify a template and the endpoint of the dynamic tunnels (that is, the destination network). The destination network could be one device or multiple devices (indicated by a subnet). The template (**odnmytemplate** in this example) is specified under the *protocols source-packet-routing source-routing-path-template*. The device configuration also points to a compute profile which can include additional parameters.

The configuration on the device establishes:

- The source routing path template
- The destination network
- The compute profile

See the [Junos OS documentation](#) for general guidelines on the syntax and usage of these specific commands.

Once the router creates the dynamic tunnels, run the `show dynamic-tunnels database` command to view the new tunnel as shown in the following example. The dynamic tunnel also appears in the Tunnel tab in the network information table.

```
northstar@PE1# run show dynamic-tunnels database
*- Signal Tunnels #- PFE-down
Table: inet.3
Destination-network: 10.0.0.11/32
Tunnel to: 10.0.0.11/32
Reference count: 1
Next-hop type: spring-te
10.0.0.11:dt-srte-odncf
State: Established
```

To delegate an ODN tunnel to the PCE, configure the following statements on the device running Junos OS:

```
set protocols source-packet-routing lsp-external-controller pccd
set protocols source-packet-routing source-routing-path-template odnmytemplate
lsp-external-controller pccd
```

NOTE: In the commands above, **odnmytmeplate** is the name of a particular template.

View the Segment Routing Path

To view the details of the segment routing path, do one of the following:

- The IP address and the SID are the two parts of an explicit route. The IP address part is displayed in the ERO column and the SID part is displayed in the Record Route column in the Tunnel tab of the network information table.
- In the Tunnel tab, hover over a tunnel and click the **Details** icon that appears. Alternatively, select a tunnel and click **More > Show Detail**.

The Tunnel-*<Tunnel Name>* page appears. Navigate to the Details tab and click **liveProperties > ero [n]** to view details of the ERO.

- Use Junos OS show commands on the router. Some examples are:
 - `show spring-traffic-engineering lsp name lsp-name detail` to display the tunnel status and SID labels.

- `show route table inet.3` to display the mapping of traffic destinations with SPRING tunnels.

Points to Remember

Following are a few things to keep in mind with regard to SRv6 tunnels:

- You can provision SRv6 tunnels only on JUNOS routers.
- Paragon Pathfinder does not support the collection of telemetry data for SRv6 tunnels.
- You must configure the End.X SID on the ISIS interfaces through which adjacency is made, for the link to be considered in the SRv6 path computation.
- The current path for SRv6 tunnels is not displayed when you view the events for such tunnels.
- You cannot run ping and traceroute commands for SRv6 tunnels.
- You can provision SRv6 tunnels only by using the default Path Computation Server.
- Paragon Pathfinder supports only one SRv6 SID per node and per link.
- You cannot provision segment routing policies on Cisco IOS-XR devices by using NETCONF as the provisioning method.

Add a Single Tunnel

You can provision tunnels by using either the Path Computation Element Protocol (PCEP) or the Network Configuration Protocol (NETCONF). Whether provisioned by using PCEP or NETCONF, tunnels can be learned by using PCEP or device collection. If learned by using device collection, the PCE requires periodic device collection to learn about tunnels and other updates to the network. To learn about how to schedule a device collection task, see ["Add a Device Collection Task" on page 978](#).

NOTE: For Cisco IOS-XR devices, you must first run device collection before provisioning tunnels by using NETCONF.

Once you create device collection tasks, the PCE discovers tunnels provisioned by using NETCONF. Unlike PCEP, the PCE with NETCONF supports logical nodes.

For more information about managing logical nodes, see *Considerations When Using Logical Nodes* later in this topic.

To provision a single tunnel:

1. Select **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. In the Tunnel tab, select **Provisioning > Tunnel**.

The Add Tunnel page appears.

3. Complete the configuration on each tab according to the guidelines in [Table 114 on page 730](#).

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

4. (Optional) From any tab, click **Preview Path** at the bottom of the page to view the path on the topology map.

5. 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:

- If you have the Auto-approve permission assigned to your user role, the request is automatically approved and deployed on the devices.
- If you don't have Auto-approve permission, the request must be manually approved and then, deployed. See ["About the Change Control Management Page" on page 829](#).

The tunnel then appears in the Tunnel tab of the network information table (in the Topology page).

Table 114: Fields on the Add Tunnel Page

Field	Description
-------	-------------

Properties

Table 114: Fields on the Add Tunnel Page (*Continued*)

Field	Description
Provisioning Method	<p>From the list, select one of the following methods to be used to provision the tunnel:</p> <ul style="list-style-type: none"> • NETCONF (default)—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. • 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. <p>NOTE:</p> <ul style="list-style-type: none"> • 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	<p>From the list, select the type of tunnel that you want to provision:</p> <ul style="list-style-type: none"> • RSVP • SR (segment routing) • SRv6
Name	<p>For a primary tunnel, specify a unique name for the tunnel.</p> <p>For a secondary or standby tunnel, specify the same name as the primary tunnel that is associated with the secondary or standby tunnel.</p> <p>You can use any number of alphanumeric characters, hyphens, and underscores.</p> <p>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.</p>
Node A	<p>From the list, select the node that you want to use as the ingress node.</p>

Table 114: Fields on the Add Tunnel Page (*Continued*)

Field	Description
Node Z	From the list, select the node that you want to use as the egress node.
IP Z	<p>From the list, select the IP address for Node Z (that is, the egress node).</p> <p>The options in the list are populated based on the Node Z that you selected.</p>
Admin Status	<p>The Path Computation Server uses the administration status of the tunnel to decide whether to route, provision, or both route and provision the tunnel.</p> <p>If the Path Computation Server routes the tunnel, no traffic flows through the tunnel and its operational status is Up. If the Path Computation Server provisions the tunnel, traffic flows through the tunnel and its operational status is Active.</p> <p>Select one of the following options as the administration status:</p> <ul style="list-style-type: none"> • Up—If you select this option, the Path Computation Server routes and provisions the tunnel. • Planned—If you select this option, the Path Computation Server routes the tunnel and reserves capacities for the tunnel. However, the Path Computation Server doesn't provision the tunnel. • Shutdown—If you select this option, the Path Computation Server neither routes nor provisions the tunnel. The tunnel is maintained in the datastore and is associated with a persist state. This means that the tunnel can be brought back up at a later time, if required.
Path Type	From the list, select primary, secondary, or standby as the path type.
Path Name	<p>Specify the name for the path.</p> <p>This field is available only for primary tunnels with RSVP provisioning type, and for all secondary and standby tunnels.</p>

Table 114: Fields on the Add Tunnel Page (*Continued*)

Field	Description
Planned Bandwidth	<p>Specify the planned bandwidth (along with valid units, with no space between the bandwidth and units) for the tunnel.</p> <p>If you specify a value without units, bps is automatically applied.</p> <p>Valid units are:</p> <ul style="list-style-type: none"> • B or b • M or m • K or k • G or g <p>Examples: 50M, 1000b, 25g.</p>
Bandwidth Sizing	<p>NOTE: This option is displayed only when you select PCEP as the provisioning method.</p> <p>Click the toggle button to enable or disable (default) bandwidth sizing for the tunnel.</p> <p>If you enable bandwidth sizing, the tunnel is included in the periodic re-computation of planned bandwidth based on aggregated tunnel traffic statistics.</p>
Adjustment Threshold (%)	<p>NOTE: This option is available only when you enable bandwidth sizing.</p> <p>Specify the sensitivity (in %) of the automatic bandwidth adjustment.</p> <p>The new planned bandwidth is only considered if it differs from the existing bandwidth by the value of this setting or more. The default value is 10%.</p>

Table 114: Fields on the Add Tunnel Page (Continued)

Field	Description
Minimum Bandwidth	<p>NOTE: This option is available only when you enable bandwidth sizing.</p> <p>Specify the minimum 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.</p> <p>If the new planned bandwidth is less than the minimum setting, the PCE signals the tunnel with the minimum bandwidth. However, if the new planned bandwidth falls in between the maximum and minimum settings, The PCE signals the tunnel with the new planned bandwidth.</p> <p>The valid units are:</p> <ul style="list-style-type: none"> • B or b • M or m • K or k • G or g <p>Examples: 50M, 1000b, 25g.</p>
Maximum Bandwidth	<p>NOTE: This option is available only when you enable bandwidth sizing.</p> <p>Specify the maximum 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.</p> <p>If the new planned bandwidth is greater than the maximum setting, the PCE signals the tunnel with the maximum bandwidth. However, if the new planned bandwidth falls in between the maximum and minimum settings, the PCE signals the tunnel with the new planned bandwidth.</p> <p>The valid units are:</p> <ul style="list-style-type: none"> • B or b • M or m • K or k • G or g <p>Examples: 50M, 1000b, 25g.</p>

Table 114: Fields on the Add Tunnel Page (Continued)

Field	Description
Minimum Variation Threshold	<p>NOTE: This option is available only when you enable bandwidth sizing.</p> <p>Specify the sensitivity of the automatic bandwidth adjustment when the new planned bandwidth is compared with the current planned bandwidth.</p> <p>Default: Zero.</p> <p>The new planned bandwidth is only considered if the difference is greater than or equal to the value of this setting. Because it is not a percentage, this is used to prevent small fluctuations from triggering unnecessary bandwidth changes.</p> <p>If both the adjustment threshold and the minimum variation threshold are greater than zero, both settings are taken into consideration. The new planned bandwidth is considered if the percentage difference is greater than or equal to the adjustment threshold and the actual difference is greater than or equal to the minimum variation.</p>
Color Community	<p>NOTE: This field is available only for segment routing tunnels.</p> <p>Assign a color for the segment routing tunnel that can be used to map traffic on the tunnel.</p>
Use Penultimate Hop as Signaling Address	<p>NOTE: This field is available only for segment routing tunnels.</p> <p>Click the toggle button to enable the Path Computation Server to use the penultimate hop as the signaling address for Egress Peer Engineering.</p> <p>If you haven't specified a color community, the setting applies to all traffic. If you've specified a color community, the setting applies to traffic in that color community.</p>
Setup	<p>Specify the setup priority for the tunnel traffic.</p> <p>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.</p> <p>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.</p>

Table 114: Fields on the Add Tunnel Page (*Continued*)

Field	Description
Hold	<p>Specify the hold priority for the tunnel traffic.</p> <p>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.</p> <p>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.</p>
Planned Metric	<p>Specify the static tunnel metric.</p> <p>The PCE uses this metric to route the tunnel instead of allowing the router to choose a path.</p>
Routing Method	<p>From the list, select a routing method to specify whether the PCE should compute and provision the path for the tunnel:</p> <p>The available options are:</p> <ul style="list-style-type: none"> • 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. <p>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.</p> <p>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.</p> <ul style="list-style-type: none"> • Other routing methods (default, delay, adminWeight, constant, distance, 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. <p>For example, a loose explicit route specified by you or learned from the router might be modified to a strict explicit route.</p>

Table 114: Fields on the Add Tunnel Page (*Continued*)

Field	Description
Binding SID	<p>NOTE: Binding SID field is available only for segment routing tunnels with NETCONF as the provisioning type.</p> <p>Specify the numerical binding SID label value.</p> <p>Binding SID represents the path that is defined by the hops you specify on the Path tab (that is, the hops that make up the private forwarding adjacency link).</p> <p>Range: 1000000 to 1048575.</p>
Use Node SID	<p>NOTE: Use Node SID field is valid only for tunnels with SR as the provision type. To use Node SID for path computation, you must configure the LSP to Path Computation Instance fields in the Pathfinder Settings. For more information, see "LSP to Path Computation Instance" on page 224.</p> <p>Enable this field to use Node SIDs for path computation.</p>
<i>Constraints</i>	
Admin Group Include All	<p>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.</p> <p>The admin group bits are mapped to meaningful names, such as colors (configured from the Configuration > Network > Admin Group page). You can easily differentiate the different traffic routes in the display and also use coloring constraints to influence the path of the tunnel.</p>
Admin Group Include Any	<p>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.</p>
Admin Group Exclude	<p>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.</p>
Maximum Delay	<p>Specify the maximum delay (in milliseconds) for the tunnel, which is used as a constraint for tunnel rerouting.</p>
Maximum Hop	<p>Specify an integer value for the maximum number of hops that the tunnel can traverse.</p>

Table 114: 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.
<i>Advanced</i>	
Count	<p>Specify the number of parallel tunnels to be created between two endpoints.</p> <p>These tunnels share the same design parameters as specified in the Constraints tab.</p> <p>NOTE: Creating parallel tunnels in this manner is different from provisioning multiple tunnels (Provisioning > Multiple Tunnels) where you configure Design parameters separately for each tunnel.</p>
Delimiter	<p>NOTE: This field is available only when the count value is greater than 1.</p> <p>Specify a delimiter value, which can consist of alphanumeric characters and special characters except space, comma (,), and semicolon (;).</p> <p>This value is used in the automatic naming of parallel tunnels that share the same design parameters. The PCE names the tunnels using the name you enter in the Properties tab and appends the delimiter value plus a unique numerical value beginning with 1.</p> <p>Example: myTunnel_1, myTunnel_2, and so on.</p>
Description	Specify a comment or description for the tunnel for your reference.
Symmetric Pair Group	<p>Specify a unique name for the symmetric pair group. You can use any number of alphanumeric and special characters.</p> <p>Tunnels with the same group name (as specified in this field) are considered part of a symmetric pair group.</p> <p>You create a symmetric pair group so that the tunnel from the ingress node to the egress node follows the same path as the tunnel from the egress node to the ingress node. When two tunnels are present with the same end nodes but in opposite directions, the path routing uses the same set of links. For example, assume that the source to destination for Tunnel1 is NodeA to NodeZ, and the source to destination for Tunnel2 is NodeZ to NodeA. Selecting Tunnel1-Tunnel2 as a symmetric pair group places both tunnels along the same set of links. Tunnels in the same group are paired based on the source and destination node.</p>

Table 114: Fields on the Add Tunnel Page (*Continued*)

Field	Description
Create Symmetric Pair	<p>NOTE: This option is displayed only when you specify a symmetric pair group.</p> <p>Click the toggle button to enable the creation of a symmetric pair.</p> <p>This option allows you to create the symmetric pair in the same operation as creating the tunnel.</p>
Diversity Group	Specify the name for a group of tunnels to which this tunnel belongs, and for which you want diverse paths.
Diversity Level	<p>From the list, select the level of diversity for the tunnel:</p> <ul style="list-style-type: none"> • Default—No diversity level will be applied. • 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. • 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. • Link—Two paths don't intersect at any given link. Link diversity is the weakest.
Slice Include All	Specify one or more topology slice IDs for the tunnel to be routed over links and nodes, that are tagged with all the slice IDs specified in this field.
Slice Include Any	Specify one or more topology slice IDs for the tunnel to be routed over links and nodes, that are tagged with at least one of the slice IDs specified in this field.
Slice Exclude	Specify one or more topology slice IDs for the tunnel to be routed over links and nodes, that are not tagged with any of the slice IDs specified in this field.
Route on Protected IP Link	Click to enable the toggle button if you want the route to use protected IP links as much as possible.

Table 114: Fields on the Add Tunnel Page (*Continued*)

Field	Description
Custom Attributes	<p>Click the Add icon (+) to specify provisioning properties not directly supported by the GUI.</p> <p>For example, you cannot specify a hop-limit when you provision a tunnel. However, you can add hop-limit as a custom attribute.</p> <p>At the edit > protocols > mpls > label-switched-path hierarchy level in the NETCONF template file, you must add the statements that are needed to provision with the property you are adding. If the property is present with the defined value, then the provisioning statement is executed.</p>
<i>Path</i>	
Routing Path Type	<p>From the list, select the type of routing path for the tunnel:</p> <ul style="list-style-type: none"> • 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 114: Fields on the Add Tunnel Page (*Continued*)

Field	Description
Add Hop	<p>This option available only if the routing path type is Preferred or Required.</p> <p>Click the Add (+) icon or click Add Hop. From the list, select an option as the first hop between node A and node Z.</p> <p>NOTE: For SRv6 provisioning type, the list displays IPv6 router identifiers.</p> <p>In addition, click the toggle button next to this field to specify whether the hop is strict or loose:</p> <ul style="list-style-type: none"> • 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. <p>To add more hops, click the + icon again. You can add a maximum of 37 hops.</p> <p>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.</p>
<i>Schedule</i>	
Plan	<ul style="list-style-type: none"> • No Schedule—(Default) tunnel provisioning is not scheduled (that is, tunnels are provisioned immediately upon submission of the provisioning request). • Once—In the Start and End fields that appear, specify the start date and time and end date and time at which you want to provision the tunnels. The tunnels are provisioned once at the specified date and time. • Recurring Daily—Specify the start and end dates and start and end times in the Start Date, End Date, Start Time, and End Time fields that appear. The tunnels are provisioned daily.

Considerations When Using Logical Nodes

You can add and provision tunnels that incorporate logical nodes. Junos OS does not support PCEP for logical nodes, but the PCE can still import logical node information using device collection. When you run a device collection task, the PCE uses the Junos OS **show configuration** command on each router to obtain both physical and logical node information. The logical node information must then be correlated with the physical node information, before provisioning tunnels that use logical nodes.

To provision a tunnel that uses logical nodes:

1. On the Topology page (**Network > Topology**), click the Node tab in the network information table and confirm that the PCEP Status is Up for all the physical nodes. For logical nodes, the PCEP Status is blank because the Path Computation Element Protocol cannot directly discover tunnels originating from a logical system.
2. Enable NETCONF for the physical nodes (if not already done):
 - a. Select **Configuration > Devices**.

The Device page appears.
 - b. Select a device and click the Edit icon.

The Edit *Device-Name* page appears.
 - c. In the **Protocols** section, select NETCONF and click the toggle button to enable NETCONF for the selected device.
 - d. Click **Save**.

NETCONF is now enabled for the selected device.

Repeat the procedure to enable NETCONF for multiple devices.
 - e. On the Topology page, click the Node tab in the network information table and confirm that the NETCONF Status is Up for these devices.
3. Create and run a device collection task to obtain the latest information.

NOTE: Run device collection before you attempt to create tunnels that incorporate logical nodes. Otherwise, the logical nodes are not available as selections for Nodes A and Z in the Add Tunnel page (**Network > Topology > Tunnel tab > Provisioning**).

To create a device collection task:

- a. Select **Settings > Network > Task Scheduler**.

The Task Scheduler page appears.
- b. Click **Add**.

The Create New Task page (wizard) appears.
- c. Configure the fields on each step of the wizard, as required.

If you use the Selective Devices option, select only the physical devices. See ["Add a Device Collection Task" on page 978](#) for more information.

d. Click **Submit**.

The details of the task that you created are displayed on the Task Scheduler page. The device collection data is sent to the Path Computation Server for routing and is reflected in the Topology view.

When you run the device collection task, the PCE uses the Junos OS **show configuration** command on each physical router to obtain both physical and logical node information. This information enables the PCE to correlate each logical node to its corresponding physical node. You can confirm this correlation from the Node tab in the network information table (**Network > Topology**).

4. (Optional) Add the Physical Hostname and Physical Host IP columns to the Node tab. For a logical node, the hostname and IP address in these columns tell you which physical node correlates with the logical node.

5. Provision tunnels:

Now that the logical nodes are in the device list and are correlated to the correct physical nodes, you can create tunnels that incorporate logical nodes. You do this using the same procedure as tunnels using only physical nodes. Ensure that you select NETCONF as the provisioning method.

6. Run the device collection task periodically to keep the logical node information updated. There are no real time updates for logical nodes.

RELATED DOCUMENTATION

[About the Tunnel Tab | 710](#)

[Edit and Delete Tunnels | 764](#)

Add Diverse Tunnels

When creating a route between two sites, you might not want to rely on a single tunnel to send traffic from one site to another. By creating a second tunnel (routing path) between the two sites, you can protect your network against failures and balance the network load.

NOTE:

- If the PCE is unable to achieve the diversity level you request, it still creates the diverse tunnel pair, using a diversity level as close as possible to the level you requested.

- By default, the PCE does not reroute a diverse tunnel pair when there is a network outage. For diverse tunnels to be rerouted, you can use the Path Optimization feature (**Network > Topology > Path Optimization**) and schedule path optimization to occur at regular intervals.
- When provisioning diverse tunnels, the PCE might return an error if the value you specified in the Site field on the Edit Node page (**Network > Topology > Node tab > Edit icon**) contains special characters. Hence, we recommend using alphanumeric characters only.

To add diverse tunnels:

1. Select 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. In the Tunnel tab, select Provisioning > Diverse tunnels.

The Add Diverse Tunnels page appears.

3. Complete the configuration on each tab of the Add Diverse Tunnels page according to the guidelines in [Table 115 on page 744](#).

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

4. (Optional) Click Preview Paths at the bottom of the page to see the paths drawn on the topology map.

5. Click Add to add the tunnels.

A confirmation message appears on the top of the page, indicating that an add tunnel request was successfully created:

- If you have the Auto-approve permission assigned to your user role, the request is automatically approved and deployed on the devices.
- If you don't have Auto-approve permission, the request must be manually approved and then, deployed. See "[About the Change Control Management Page](#)" on page 829.

The tunnels then appear in the Tunnel tab of the network information table (in the Topology page).

Table 115: Fields on the Add Diverse Tunnels Page

Field	Description
-------	-------------

Properties

Table 115: Fields on the Add Diverse Tunnels Page (*Continued*)

Field	Description
Provisioning Method	<p>From the list, select one of the following methods to be used to provision the tunnel:</p> <ul style="list-style-type: none"> • NETCONF (default)—The tunnels are statically provisioned and the associated configuration statements appear in the router configuration file. Upon provisioning, these tunnels are added as device-controlled tunnels. • PCEP (Path Computation Element Protocol)—The tunnels are initiated by the path computation element (PCE) and the associated configuration statements do not appear in the router configuration file. Upon provisioning, these tunnels are added as PCE-initiated tunnels. <p>NOTE: For IOS-XR routers, NETCONF-based tunnel provisioning has the same capabilities as PCEP-based tunnel provisioning.</p>
Provisioning Type	<p>From the list, select the type of tunnel that you want to provision:</p> <ul style="list-style-type: none"> • RSVP • Segment Routing (SR) • SRv6
Diversity Group	<p>Specify the name of a group of tunnels to which this tunnel belongs, and for which diverse paths are desired.</p>

Table 115: Fields on the Add Diverse Tunnels Page (Continued)

Field	Description
Diversity Level	<p>From the list, select the level of diversity for the tunnel:</p> <ul style="list-style-type: none"> • Default—No diversity level is applied. • 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. • 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. <p>NOTE: If two paths are SRLG-diverse, the paths will not be routed over links which are in the same SRLG. The SRLGs can be learned from network protocols such as BGP-LS and NETCONF. The SRLGs can change dynamically and are displayed in the SRLG/Facility tab in the network information table. If the links become part of the same SRLG, the previously computed tunnels are no more diverse. In this scenario, new tunnels will not be automatically computed and applied in the network. You must manually select each tunnel (from the Tunnel tab) that is no longer diverse and request reprovisioning (Tunnel tab > Provisioning > Reprovision).</p> <ul style="list-style-type: none"> • Link—Two paths don't intersect at any given link. Link diversity is the weakest.
Tunnel 1	
NOTE: The same fields are available for Tunnel 2. Use the same guidelines to configure the fields for Tunnel 2.	
Name	Specify a unique name for the first tunnel. You can use any number of alphanumeric characters, hyphens, and underscores.
Node A	From the list, select the node that you want to use as the ingress node.
Node Z	From the list, select the node that you want to use as the egress node.
IP Z	<p>From the list, select the IP address for Node Z (that is, the egress node).</p> <p>The options in the list are populated based on the Node Z that you selected.</p>

Table 115: Fields on the Add Diverse Tunnels Page (Continued)

Field	Description
Admin Status	<p>The PCS uses the administration status of the tunnel to decide whether to route or provision, or both route and provision the tunnel.</p> <p>If the tunnel is routed, no traffic flows through the tunnel and its operational status is Up. If the tunnel is provisioned, traffic flows through the tunnel and its operational status is Active.</p> <p>Select one of the following options as the administration status:</p> <ul style="list-style-type: none"> • Up—If you select this option, the PCS routes and provisions the tunnel. • Planned—If you select this option, the PCS routes the tunnel and reserves capacities for the tunnel. However, the PCS doesn't provision the tunnel. • Shutdown—If you select this option, the PCS neither routes nor provisions the tunnel. The tunnel is maintained in the datastore and is associated with a persist state so that the tunnel can be brought back up at a later time, if required.
Planned Bandwidth	<p>Specify the planned bandwidth (along with valid units, with no space between the bandwidth and units) for the tunnel.</p> <p>If you specify a value without units, bps is applied.</p> <p>The valid units are:</p> <ul style="list-style-type: none"> • B or b • M or m • K or k • G or g <p>Examples: 50M, 1000b, 25g.</p>
Planned Metric	<p>Specify the static tunnel metric.</p> <p>The PCE uses this metric to route the tunnel instead of allowing the router itself to choose a path.</p>

Table 115: Fields on the Add Diverse Tunnels Page (Continued)

Field	Description
Bandwidth Sizing	<p>NOTE: This option is displayed only when you select PCEP as the provisioning method.</p> <p>Click the toggle button to enable bandwidth sizing for the tunnel.</p> <p>If enabled, the tunnel is included in the periodic re-computation of planned bandwidth based on aggregated tunnel traffic statistics.</p> <p>If you enable bandwidth sizing, you must configure the following parameters:</p> <ul style="list-style-type: none"> • Adjustment Threshold • Minimum Bandwidth • Maximum Bandwidth • Minimum Variation Threshold
Adjustment Threshold (%)	<p>NOTE: This option is available only when you enable bandwidth sizing.</p> <p>Specify the sensitivity (in %) of the automatic bandwidth adjustment.</p> <p>The new planned bandwidth is considered only if it differs from the existing bandwidth by the value of this setting or more. The default value is 10%.</p>

Table 115: Fields on the Add Diverse Tunnels Page (*Continued*)

Field	Description
Minimum Bandwidth	<p>NOTE: This option is available only when you enable bandwidth sizing.</p> <p>Specify the minimum 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 applied.</p> <p>If the new planned bandwidth is less than the minimum setting, the PCE signals the tunnel with the minimum bandwidth. However, if the new planned bandwidth falls in between the maximum and minimum settings, the PCE signals the tunnel with the new planned bandwidth.</p> <p>The valid units are:</p> <ul style="list-style-type: none"> • B or b • M or m • K or k • G or g <p>Examples: 50M, 1000b, 25g.</p>

Table 115: Fields on the Add Diverse Tunnels Page (Continued)

Field	Description
Maximum Bandwidth	<p>NOTE: This option is available only when you enable bandwidth sizing.</p> <p>Specify the maximum 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 applied.</p> <p>If the new planned bandwidth is greater than the maximum setting, the PCE signals the tunnel with the maximum bandwidth. However, if the new planned bandwidth falls in between the maximum and minimum settings, the PCE signals the tunnel with the new planned bandwidth.</p> <p>The valid units are:</p> <ul style="list-style-type: none"> • B or b • M or m • K or k • G or g <p>Examples: 50M, 1000b, 25g.</p>
Minimum Variation Threshold	<p>NOTE: This option is available only when you enable bandwidth sizing.</p> <p>Specify the sensitivity of the automatic bandwidth adjustment when the new planned bandwidth is compared to the current planned bandwidth.</p> <p>Default: Zero.</p> <p>The new planned bandwidth is only considered if the difference is greater than or equal to the value of this setting. Because it is not a percentage, this can be used to prevent small fluctuations from triggering unnecessary bandwidth changes.</p> <p>If both the adjustment threshold and the minimum variation threshold are greater than zero, both settings are taken into consideration. In that case, the new planned bandwidth is considered if the percentage difference is greater than or equal to the adjustment threshold, and, the actual difference is greater than or equal to the minimum variation.</p>

Table 115: Fields on the Add Diverse Tunnels Page (Continued)

Field	Description
Color Community	<p>NOTE: This field is available only for segment routing tunnels.</p> <p>Assign a color for the segment routing tunnel that can be used to map traffic on the tunnel.</p>
Use Penultimate Hop as Signaling Address	<p>NOTE: This field is available only for segment routing tunnels.</p> <p>Click the toggle button to enable the PCS to use the penultimate hop as the signaling address for Egress Peer Engineering (EPE).</p> <p>If you haven't specified a color community, the setting applies to all traffic. If you've specified a color community, the setting applies to traffic in that color community.</p>
Setup	<p>Specify the setup priority for the tunnel traffic.</p> <p>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.</p> <p>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.</p>
Hold	<p>Specify the hold priority for the tunnel traffic.</p> <p>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.</p> <p>Based on the hold priority, the PCE determines whether the tunnel can be preempted or not. If the hold priority for an tunnel is higher, it is unlikely for the tunnel to be preempted.</p>

Table 115: Fields on the Add Diverse Tunnels Page (Continued)

Field	Description
Routing Method	<p>From the list, select a routing method for the tunnel to specify whether the PCE should compute and provision the path for the tunnel, or not :</p> <p>The available options are:</p> <ul style="list-style-type: none"> • routeByDevice—This is the default routing method when a PCC-controlled tunnel is created or learned by the PCE. For this method, the PCE does not compute and provision a path. <p>This method is appropriate for three types of tunnels: RSVP TE PCC-controlled tunnels, segment routing Path Computation Element Protocol-based tunnels, and segment routing NETCONF-based tunnels.</p> <ul style="list-style-type: none"> • Other routing methods (default, delay, adminWeight, constant, distance, ISIS, OSPF)—When a PCC-controlled tunnel has a routing method that is not routeByDevice, the PCE computes and provisions the path as a strict explicit route when provisioning the tunnel. The tunnel's existing explicit route might be modified to a the PCE-computed strict explicit route. For example, a loose explicit route specified by you or learned from the router would be modified to a strict explicit route.

Constraints

NOTE: The same fields are available for Tunnel 2. Use the same guidelines to configure the fields for Tunnel 2.

Tunnel 1

Admin Group Include All	<p>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. You can select a maximum of 32 admin group bits.</p> <p>The admin group bits are mapped to meaningful names, such as colors (configured in the Configuration > Network > Admin Group page), so that you can easily differentiate the different traffic routes in the display and also use coloring constraints to influence the path of the tunnel.</p>
Admin Group Include Any	<p>From the list, select one or more admin group bits for the tunnel to traverse links that include at least one of the admin groups specified in this field. The maximum selections allowed is 32.</p>

Table 115: Fields on the Add Diverse Tunnels Page (Continued)

Field	Description
Admin Group Exclude	From the list, select one or more admin group bits for the tunnel to traverse links that do not include any of the admin groups specified in this field. The maximum selections allowed is 32.

Advanced

NOTE: The same fields are available for Tunnel 2. Use the same guidelines to configure the fields for Tunnel 2.

Tunnel 1

Description	Specify a comment or description for the tunnel for your reference.
Symmetric Pair Group	<p>Specify a unique name for the symmetric pair group. You can use any number of alphanumeric and special characters.</p> <p>tunnels with the same group name (as specified in this field) are considered part of a symmetric pair group.</p> <p>You create a symmetric pair group so that the tunnel from the ingress node to the egress node follows the same path as the tunnel from the egress node to the ingress node. When there are two tunnels with the same end nodes but in opposite directions, the path routing uses the same set of links. For example, suppose Tunnel1 source to destination is NodeA to NodeZ, and Tunnel2 source to destination is NodeZ to NodeA. Selecting Tunnel1-Tunnel2 as a symmetric pair group places both tunnels along the same set of links. Tunnels in the same group are paired based on the source and destination node.</p>
Create Symmetric Pair	<p>NOTE: This option is displayed only when you specify a symmetric pair group.</p> <p>Click the toggle button to enable the creation of a symmetric pair.</p> <p>This option allows you to create the symmetric pair in the same operation as creating the diverse tunnel.</p>

Table 115: Fields on the Add Diverse Tunnels Page (Continued)

Field	Description
Custom Attributes	<p>Click the Add icon (+) to specify provisioning properties not directly supported by the GUI.</p> <p>For example, you cannot specify a hop-limit when you provision a tunnel. However, you can add hop-limit as a custom attribute.</p> <p>At the edit > protocols > mpls > label-switched-path hierarchy level in the NETCONF template file, you must add the statements needed to provision with the property you are adding. If the property is present with the defined value, then the provisioning statement is executed.</p>
<i>Schedule</i>	
Plan	<ul style="list-style-type: none"> • No Schedule—(Default) tunnel provisioning is not scheduled (that is, tunnels are provisioned immediately upon submission of the provisioning request). • Once—In the Start and End fields that appear, specify the start date and time and end date and time at which you want to provision the tunnels. The tunnels are provisioned once at the specified date and time. • Recurring Daily—Specify the start and end dates and start and end times in the Start Date, End Date, Start Time, and End Time fields that appear. The tunnels are provisioned daily.

RELATED DOCUMENTATION

[About the Tunnel Tab | 710](#)

[Edit and Delete Tunnels | 764](#)

Add Multiple Tunnels

You can provision multiple tunnels at once in the network topology from the Add Multiple Tunnels page.

To provision multiple tunnels:

1. Select **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. In the Tunnel tab, select **Provisioning > Multiple Tunnels**.

The Add Multiple Tunnels page appears.

3. Complete the configuration on each tab according to the guidelines in [Table 116 on page 755](#).

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

4. Click **Add** to add the tunnels.

A confirmation message appears on the top of the page, indicating that an add tunnel request was successfully created:

- If you have the Auto-approve permission assigned to your user role, the request is automatically approved and deployed on the devices.
- If you don't have Auto-approve permission, the request must be manually approved and then, deployed. See "[About the Change Control Management Page](#)" on page 829.

The tunnels then appear in the Tunnel tab of the network information table (in the Topology page).

Table 116: Fields on the Add Multiple Tunnels Page

Field	Description
<i>Properties</i>	
Provisioning Method	<p>From the list, select one of the following methods to be used to provision the tunnel:</p> <ul style="list-style-type: none"> • NETCONF (default)—The tunnels are statically provisioned and the associated configuration statements appear in the router configuration file. Upon provisioning, these tunnels are added as PCC-controlled tunnels or device-controlled tunnels. • PCEP (Path Computation Element Protocol)—The tunnels are initiated by the path computation element (PCE) and the associated configuration statements do not appear in the router configuration file. Upon provisioning, these tunnels are added as PCE-initiated tunnels. <p>NOTE:</p> <ul style="list-style-type: none"> • For IOS-XR routers, NETCONF-based tunnel provisioning has the same capabilities as Path Computation Element Protocol-based tunnel provisioning.

Table 116: Fields on the Add Multiple Tunnels Page (Continued)

Field	Description
Provisioning Type	<p>From the list, select the type of tunnel that you want to provision:</p> <ul style="list-style-type: none"> • RSVP • Segment Routing (SR) • SRv6
ID Prefix	<p>Specify a prefix to be applied to the names of all the tunnels that are created.</p> <p>Default: PCE.</p>
Node Z Tag	<p>From the list, select a tag as the secondary loopback address for Node Z.</p> <p>The list is populated from the tags that you specify in the Advanced tab of the Modify Node page (Network > Topology > Node tab > Edit icon), where you add destination IP addresses in addition to the default IPv4 router ID address, and assign a descriptive tag to each.</p>
Node A List	<p>Select one or more nodes to be part of the Node A list.</p> <p>For a full mesh tunnel to be created, you can specify the same nodes for Node A and Node Z by clicking the <i>Copy Node Z List</i> link (that is located above the list). All the nodes that you specify in the Node Z List are added to the Node A List.</p>
Node Z List	<p>Select one or more nodes to be part of the Node Z list.</p> <p>For a full mesh tunnel to be created, you can specify the same nodes for Node Z and Node A by clicking the <i>Copy Node A List</i> link (that is located above the list). All the nodes that you specified in the Node A List are added to the Node Z List.</p>

Table 116: Fields on the Add Multiple Tunnels Page (Continued)

Field	Description
Admin Status	<p>The PCS uses the administration status of the tunnel to decide whether to route or provision, or both route and provision the tunnel.</p> <p>If the tunnel is routed, no traffic flows through the tunnel and its operational status is Up. If the tunnel is provisioned, traffic flows through the tunnel and its operational status is Active.</p> <p>Select one of the following options as the administration status:</p> <ul style="list-style-type: none"> • Up—If you select this option, the PCS routes and provisions the tunnel. • Planned—If you select this option, the PCS routes the tunnel and reserves capacities for the tunnel. However, the PCS doesn't provision the tunnel. • Shutdown—If you select this option, the PCS neither routes nor provisions the tunnel. The tunnel is maintained in the datastore and is associated with a persist state so that the tunnel can be brought back up at a later time, if required.
Planned Bandwidth	<p>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 applied.</p> <p>Valid units are:</p> <ul style="list-style-type: none"> • B or b • M or m • K or k • G or g <p>Examples: 50M, 1000b, 25g.</p>

Table 116: Fields on the Add Multiple Tunnels Page (Continued)

Field	Description
Bandwidth Sizing	<p>NOTE: This option is displayed only when you select Path Computation Element Protocol as the provisioning method.</p> <p>Click the toggle button to enable bandwidth sizing for the tunnel.</p> <p>If enabled, the tunnel is included in the periodic re-computation of planned bandwidth based on aggregated tunnel traffic statistics.</p> <p>If you enable bandwidth sizing, you must configure the following parameters:</p> <ul style="list-style-type: none"> • Adjustment Threshold • Minimum Bandwidth • Maximum Bandwidth • Minimum Variation Threshold
Adjustment Threshold (%)	<p>NOTE: This option is available only when you enable bandwidth sizing.</p> <p>Specify the sensitivity (in %) of the automatic bandwidth adjustment.</p> <p>The new planned bandwidth is considered only if it differs from the existing bandwidth by the value of this setting or more. The default value is 10%.</p>

Table 116: Fields on the Add Multiple Tunnels Page (*Continued*)

Field	Description
Minimum Bandwidth	<p>NOTE: This option is available only when you enable bandwidth sizing.</p> <p>Specify the minimum 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 applied.</p> <p>If the new planned bandwidth is less than the minimum setting, the PCE signals the tunnel with the minimum bandwidth. However, if the new planned bandwidth falls in between the maximum and minimum settings, the PCE signals the tunnel with the new planned bandwidth.</p> <p>The valid units are:</p> <ul style="list-style-type: none"> • B or b • M or m • K or k • G or g <p>Examples: 50M, 1000b, 25g.</p>
Maximum Bandwidth	<p>NOTE: This option is available only when you enable bandwidth sizing.</p> <p>Specify the maximum 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 applied.</p> <p>If the new planned bandwidth is greater than the maximum setting, the PCE signals the tunnel with the maximum bandwidth. However, if the new planned bandwidth falls in between the maximum and minimum settings, the PCE signals the tunnel with the new planned bandwidth.</p> <p>The valid units are:</p> <ul style="list-style-type: none"> • B or b • M or m • K or k • G or g <p>Examples: 50M, 1000b, 25g.</p>

Table 116: Fields on the Add Multiple Tunnels Page (Continued)

Field	Description
Minimum Variation Threshold	<p>NOTE: This option is available only when you enable bandwidth sizing.</p> <p>Specify the sensitivity of the automatic bandwidth adjustment when the new planned bandwidth is compared to the current planned bandwidth.</p> <p>Default: Zero.</p> <p>The new planned bandwidth is only considered if the difference is greater than or equal to the value of this setting. Because it is not a percentage, this can be used to prevent small fluctuations from triggering unnecessary bandwidth changes.</p> <p>If both the adjustment threshold and the minimum variation threshold are greater than zero, both settings are taken into consideration. In that case, the new planned bandwidth is considered if the percentage difference is greater than or equal to the adjustment threshold, and, the actual difference is greater than or equal to the minimum variation.</p>
Color Community	<p>NOTE: This field is available only for segment routing tunnels.</p> <p>Assign a color for the segment routing tunnel that can be used to map traffic on the tunnel.</p>
Use Penultimate Hop as Signaling Address	<p>NOTE: This field is available only for segment routing tunnels.</p> <p>Click the toggle button to enable the PCS to use the penultimate hop as the signaling address for Egress Peer Engineering (EPE).</p> <p>If you haven't specified a color community, the setting applies to all traffic. If you've specified a color community, the setting applies to traffic in that color community.</p>
Setup	<p>Specify the setup priority for the tunnel traffic.</p> <p>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.</p> <p>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.</p>

Table 116: Fields on the Add Multiple Tunnels Page (*Continued*)

Field	Description
Hold	<p>Specify the hold priority for the tunnel traffic.</p> <p>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.</p> <p>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.</p>
Planned Metric	<p>Specify the static tunnel metric.</p> <p>The PCE uses this metric to route the tunnel instead of allowing the router itself to choose a path.</p>
Routing Method	<p>From the list, select a routing method for the tunnel to specify whether the PCE should compute and provision the path for the tunnel, or not:</p> <p>The available options are:</p> <ul style="list-style-type: none"> • routeByDevice—This is the default routing method when a PCC-controlled tunnel is created or learned by the PCE. For this method, the PCE does not compute and provision a path. <p>This method is appropriate for three types of tunnels: RSVP TE PCC-controlled tunnels, Segment routing Path Computation Element Protocol-based tunnels, and Segment routing NETCONF-based tunnels.</p> • Other routing methods (default, delay, adminWeight, constant, distance, ISIS, OSPF)—When a PCC-controlled tunnel has a routing method that is not routeByDevice, the PCE computes and provisions the path as a strict explicit route when provisioning the tunnel. 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 would be modified to a strict explicit route.

Constraints

Table 116: Fields on the Add Multiple Tunnels Page (Continued)

Field	Description
Admin Group Include All	<p>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. You can select a maximum of 32 admin group bits.</p> <p>The admin group bits are mapped to meaningful names (such as colors) on the Admin Group page (Configuration > Network > Admin Group). This enables you to easily differentiate the different traffic routes in the display and also use coloring constraints to influence the path of the tunnel.</p>
Admin Group Include Any	From the list, select one or more admin group bits for the tunnel to traverse 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 for the tunnel to traverse links that do not include any of the admin groups specified in this field. The maximum selections allowed is 32.
<i>Advanced</i>	
Count	<p>Specify the number of copies of the tunnels to be created (Default: 1).</p> <p>Example: if you specify a count of 2, two copies of each tunnel are created.</p>
Delimiter	<p>NOTE: This field is available only when the Count value is greater than 1.</p> <p>Specify a delimiter value, which can consist of alphanumeric characters and special characters except space, comma (,) and semicolon (;).</p> <p>This value is used in the automatic naming of parallel tunnels that share the same design parameters. The PCE names the tunnels using the name you enter in the Properties tab and appends the delimiter value plus a unique numerical value beginning with 1</p> <p>Example: mytunnel_1, mytunnel_2, and so on.</p>
Description	Specify a comment or description for the tunnel for your reference.
Diversity Group	Specify the name of a group of tunnels to which this tunnel belongs, and for which diverse paths are desired.

Table 116: Fields on the Add Multiple Tunnels Page (Continued)

Field	Description
Diversity Level	<p>From the list, select the level of diversity for the tunnel:</p> <ul style="list-style-type: none"> • Default—No diversity level is applied. • 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. • 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. • Link—Two paths don't intersect at any given link. Link diversity is the weakest.
Custom Attributes	<p>Click the Add icon (+) to specify provisioning properties not directly supported by the GUI.</p> <p>For example, you cannot specify a hop-limit when you provision a tunnel. However, you can add hop-limit as a custom attribute.</p> <p>At the edit > protocols > mpls > label-switched-path hierarchy level in the NETCONF template file, you must add the statements needed to provision with the property you are adding. If the property is present with the defined value, then the provisioning statement is executed.</p>
<i>Schedule</i>	
Plan	<p>Select one of the following plans to schedule tunnel provisioning:</p> <p>NOTE: The time zone is the server time zone.</p> <ul style="list-style-type: none"> • No Schedule—(Default) tunnel provisioning is not scheduled (that is, tunnels are provisioned immediately upon submission of the provisioning request). • Once—In the Start and End fields that appear, specify the start date and time and end date and time at which you want to provision the tunnels. The tunnels are provisioned once at the specified date and time. • Recurring Daily—Specify the start and end dates and start and end times in the Start Date, End Date, Start Time, and End Time fields that appear. The tunnels are provisioned daily.

RELATED DOCUMENTATION

[About the Tunnel Tab | 710](#)

[Edit and Delete Tunnels | 764](#)

Edit and Delete Tunnels

IN THIS SECTION

- [Edit Tunnels | 764](#)
- [Delete Tunnels | 765](#)

You can edit the parameters configured for tunnels and delete tunnels which you no longer need.

Edit Tunnels

To edit the parameters configured for a tunnel:

1. Select **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. In the Tunnel tab, select the check box corresponding to the tunnel for which you want to edit the parameters, and click the Edit icon.

The Edit Tunnel page appears, displaying the same fields that are presented when you add a tunnel.

3. Modify the parameters as needed. Based on the type of tunnel that you selected, see "[Add a Single Tunnel](#)" on page 729, "[Add Diverse Tunnels](#)" on page 743, or "[Add Multiple Tunnels](#)" on page 754 for more information on these parameters.

NOTE: You can modify only some fields when you are editing tunnel.

4. Click **Edit** to save your changes.

You are returned to the Topology page, where a confirmation message appears on the top of the page, indicating that an edit tunnel change request was successfully created:

- If you have the Auto-approve permission assigned to your user role, the request is automatically approved and deployed on the devices.

- If you don't have Auto-approve permission, the request must be manually approved and then, deployed. See "[About the Change Control Management Page](#)" on page 829.

The tunnel parameters are then updated.

Delete Tunnels

To delete one or more tunnels:

1. Select **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. In the Tunnel tab, select the check box corresponding to the tunnels that you want to delete, and click the Delete icon.

An alert message appears, asking you to confirm the delete operation.

3. Click **OK**.

You are returned to the Topology page, where a confirmation message appears, indicating that the delete tunnel change request was successfully created:

- If you have the Auto-approve permission assigned to your user role, the request is automatically approved and deployed on the devices.
- If you don't have the Auto-approve permission, the request must be manually approved and then, deployed. See "[About the Change Control Management Page](#)" on page 829.

The selected tunnels are then deleted.

RELATED DOCUMENTATION

| [About the Tunnel Tab](#) | 710

About the Demand Tab

IN THIS SECTION

- [Tasks You Can Perform](#) | 766

Use the Demand tab to view aggregated demands that are generated based on the flow monitored by the NetFlow collector.

To access this tab, select **Network > Topology** and hover over the **More** list (next to the Tunnel tab) and select **Demand**.

Tasks You Can Perform

- View the routing path for a demand—Select a demand in the network information table for the corresponding routing path to be highlighted in the topology map.

NOTE: Currently, you can preview paths on the topology map only for RSVP LSPs (not for segment-routed LSPs).

- View demand traffic—To view demand traffic over a period of time in graphical form, select a demand and click **View Demand Traffic**. In the Demand Traffic page that appears, select the period for which you want to view the data. You can view data for the past 5 minutes, 10 minutes, 1 hour, 3 hours, the previous day, the previous week, or choose from a custom time range by specifying the start and end dates (in MM/DD/YY format) and times (in HH:MM AM/PM 12-hour format) in the Choose Time Range page that appears.
You can click the refresh icon so that the page refreshes to display the most recent data. Alternatively, you can enable the **Auto Refresh** toggle button and select the interval (10 seconds, 30 seconds, 1 minute, or 5 minutes) after which the page must refresh automatically to display the latest data.
- Download details of all the demands—To view detailed information of all the demands displayed in the network information table, click **Download**. You can choose to open the comma-separated values (CSV) file with Excel or other applications, or save the file to your local system.
- From the More list, you can perform the following tasks:
 - View details of a demand—To view details of a demand, select the demand and click **Show Detail**. Alternatively, hover over a demand and click the details icon that appears. You can also right-click on a demand and select **Show Detail**.
The *Demand-Demand Name* page appears, displaying the details, such as live properties, planned properties, and operational status of the demand.
 - Provision multiple tunnels—To add multiple tunnels pertaining to a demand, select a demand and click **Provision Multiple Tunnels**. Alternatively, right-click on a demand and select **Provision Multiple Tunnels**.
The Provision Multiple Tunnels page appears, with some parameters pertaining to traffic flow (such as Node A and Node Z) pre-populated from the selected demand. You can modify the parameters as required, and provision the tunnels. See ["Add Multiple Tunnels" on page 754](#).

- Reload the network information table—To download updated demand data to your Web browser, click **Reload**. A REST API query is sent to the Paragon Automation server and then, the network information table is updated.
- Delete one or more demands—Even when a flow is no longer observed, the demand is retained and displayed in the Demand tab until you delete it. To delete such demands manually, select one or more demands and click the delete (trash can) icon. The selected demands are deleted.
Alternatively, you can add a Demand Aging task to automate the deletion. See "[Add a Demand Aging Task](#)" on page 983.

RELATED DOCUMENTATION

[NetFlow Collector Overview](#) | 864

About the Interface Tab

IN THIS SECTION

- [Tasks You Can Perform](#) | 767

You can view detailed information about interfaces associated with different nodes on the topology map from the **Interfaces** tab.

Interfaces cannot be added, modified, or deleted from the network information table.

Tasks You Can Perform

You can perform the following actions:

- From the **View** list, you can view the interface traffic (in bps) and interface delay for previous 3 hours, 1 day, 1 week, or custom time range in a graphical format. You can pin these windows anywhere on the screen using the pin icon.
- From the **Diagnostics** list, you can:

- **Show Interface**—Runs the **show interfaces** CLI command for each interface associated with the link. Interface related information is displayed on the **Diagnostics** window. You can view the status, type, node for which the command is run (From), description, and time of last execution for each associated interface. Once the commands are run, you can view detailed execution details under the **Result** tab.

NOTE: From the **New** list, you can run a new ping and traceroute command. You can also download all the diagnostic related information as a text file by clicking **Download**.

- **Download node information**—Click **Download** to download detailed information about all the existing interfaces in the topology in CSV format.
- From the **More** List, you can:
 - View detailed information about the interface by clicking **Show Detail** or click the Details icon displayed next to the interface name when you hover over it. A pop-up appears displaying the traffic and protocol details about the selected interface.
 - **Reload the network**—When there is a large number of interfaces, you can click **Reload** to clear the interface back end cache and reload the data (table entries) from the database.

Container LSP Overview

A container LSP is a logical grouping of sub-LSPs that share the properties defined in the container. Container LSPs automatically add or remove sub-LSPs based on traffic statistics. This mitigates the difficulty of finding a single path that is large enough to accommodate a large bandwidth reservation.

Using container LSPs involves:

- Adding a container LSP from the network information table (**Container LSP** tab) on the Topology page.
- Adding a container normalization task by using the Task Scheduler.
- Viewing container LSPs and their sub-LSPs in the network information table (**Container LSP** tab) on the Topology page.

RELATED DOCUMENTATION

| [Add a Container LSP](#) | 770

About the Container LSP Tab

IN THIS SECTION

- [Tasks You Can Perform | 769](#)

You can view detailed information, add, edit, or delete container LSPs from the **Container LSP** tab of the network information table on the Topology (**Network** > **Topology**) page.

Tasks You Can Perform

You can perform the following tasks:

- Add a Container LSP. See ["Add a Container LSP" on page 770](#).
- Edit the parameters of an existing Container LSP. See ["Edit Container LSP Parameters" on page 775](#).
- Delete an existing Container LSP—On the **Container LSP** tab of the network information table on the **Topology** page, select the Container LSP that you want to delete and click the delete (trash can) icon. A confirmation message appears. Click **Yes**.

The Container LSP is deleted from the network information table.

- View sub-LSPs—From the **More** list, select **Sub LSPs** to view the associated sub LSPs. You are redirected to the **Tunnel** tab where the sub LSP details are displayed. You can perform add, edit, or delete LSP functions from there.
- Download Container LSPs information—Click **Download** to download information about all the Container LSPs in CSV format.
- View Details—View detailed information about the Container LSPs by clicking **More** > **Show Detail** or click the Details icon displayed next to the Container LSP name when you hover over it.

RELATED DOCUMENTATION

[About the Topology Page | 677](#)

[Interactive Map Features Overview | 660](#)

Add a Container LSP

To add a container LSP:

1. Click **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. Hover over the **More Tabs list and select **Container LSP**.**

The Container LSP tab appears.

3. Click the add (+) icon.

The Add Container LSP page appears on the right.

4. Configure the Container LSP as per [Table 117 on page 770](#).

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

5. Click **Add.**

A confirmation message appears stating that the Add Container LSP request is successfully sent. The new Container LSP is displayed under the **Container LSP** tab.

Table 117: Fields on the Add Container LSP Page

Field	Description
Properties	
Provisioning Method	Only PCEP is available.
Provision Type	Select RSVP or SR.
Container Name	Enter name for the Container LSP. Only alphanumeric characters, hyphens, and underscores are allowed. Other special characters and spaces are not allowed. NOTE: The name you assign to the container LSP is used for automatic naming of the sub-LSPs that are created.
Node A	Click the Node A list and select the name of the ingress (source) node from the list.

Table 117: Fields on the Add Container LSP Page (Continued)

Field	Description
Node Z	Click the Node Z list and select an egress (destination) node from the list.
Merging Bandwidth	<p>Enter merging bandwidth value.</p> <p>Aggregate bandwidth thresholds are used to trigger merging of sub-LSPs during normalization. When the aggregate bandwidth usage falls below the merging bandwidth (the lower threshold), the number of sub-LSPs during normalization is reduced.</p>
Splitting Bandwidth	<p>Enter splitting bandwidth value.</p> <p>Aggregate bandwidth thresholds are used to trigger splitting of sub-LSPs during normalization. When the aggregate bandwidth usage rises above the splitting bandwidth (the upper threshold), sub-LSPs are added during normalization.</p>
Minimum and Maximum LSP Count	Enter the minimum and maximum number of sub-LSPs that can be created in the container LSP. Range is 1 through 32767.
Minimum and Maximum LSP Bandwidth	<p>Enter the minimum and maximum bandwidth that can be signaled for the sub-LSPs during normalization or initialization, immediately followed by units (no space in between). Valid units are:</p> <ul style="list-style-type: none"> • B or b (bps) • M or m (Mbps) • K or k (Kbps) • G or g (Gbps) <p>For example, 50M, 1000b, 25g.</p> <p>If you enter a value without units, bps is applied.</p>
Planned Metric	Enter the metric value for static tunnel or use the up and down arrows to increment or decrement the value. For more information, see Basic LSP Configuration in <i>MPLS Applications User Guide</i> .

Table 117: Fields on the Add Container LSP Page (Continued)

Field	Description
Setup	Enter the RSVP 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 LSP definition in Junos OS.
Hold	Enter the RSVP 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 LSP definition in Junos OS.
Routing Method	<p>Select a routing method from the following:</p> <ul style="list-style-type: none"> • Admin Weight • Constant • Default <p>NOTE: Do not change the routing method for PCEP-provisioned sub-LSPs; they should always have a routing method of “default”.</p> <ul style="list-style-type: none"> • Delay • Distance • ISIS • OSPF • Route by Device

Table 117: Fields on the Add Container LSP Page (Continued)

Field	Description
Bandwidth Sizing	<p>Toggle the button to enable bandwidth sizing. By default, this option is disabled.</p> <p>Once enabled, configure the following parameters:</p> <ul style="list-style-type: none"> • Adjustment threshold (in percentage)—Controls the sensitivity of the automatic bandwidth adjustment. The new planned bandwidth is only considered if it differs from the existing bandwidth by the value of this setting percentage or more. • Minimum and Maximum (planned) Bandwidth: <p>If the new planned bandwidth is greater than the maximum setting, Paragon PathFinder signals the LSP with the maximum bandwidth.</p> <p>If the new planned bandwidth is less than the minimum setting, Paragon PathFinder signals the LSP with the minimum bandwidth.</p> <p>If the new planned bandwidth falls in between the maximum and minimum settings, Paragon Pathfinder signals the LSP with the new planned bandwidth.</p> • Minimum Variation Threshold—Specifies the sensitivity of the automatic bandwidth adjustment when the new planned bandwidth is compared to the current planned bandwidth. The new planned bandwidth is only considered if the difference is greater than or equal to the value of this setting. This can be used to prevent small fluctuations from triggering unnecessary bandwidth changes. If both the adjustment threshold and the minimum variation threshold are greater than zero, both settings are taken into consideration. In that case, the new planned bandwidth is considered if: <ul style="list-style-type: none"> • The percentage difference is greater than or equal to the adjustment threshold and • The actual difference is greater than or equal to the minimum variation.
<p>Constraints</p>	
<p>Enables you to configure administrative groups to bypass label-switched paths (LSPs). For more information, see "Assign Names to Admin Group Bits" on page 197.</p>	
Admin Group Include All	<p>Specifies the administrative groups whose links the bypass LSP must traverse. Select one or more bit-level link coloring options from the list.</p>

Table 117: Fields on the Add Container LSP Page (Continued)

Field	Description
Admin Group Include Any	Specifies the administrative groups whose links the bypass LSP can traverse. Select one or more bit-level link coloring options from the list.
Admin Group Exclude	Specifies the administrative groups to exclude for a bypass LSP. Select one or more bit-level link coloring options from the list.

Advanced

Description	Enter a description for the Container LSP.
IP Z	IP address of Node Z.
Custom Attributes	<p>Simple name-value pairs which can be used to add any arbitrary customer-specific information. For example, to differentiate between properties for different vendor nodes.</p> <p>To add custom properties associated with the node:</p> <ol style="list-style-type: none"> 1. Click add (+) icon to add a new row. 2. Click the newly added row to enter the Name and Value. 3. Click ✓ icon to save your changes. <p>NOTE:</p> <ul style="list-style-type: none"> • You can add multiple rows (properties). • To delete an entry (row), select the row and click the delete (trash can) icon.

RELATED DOCUMENTATION

| [About the Container LSP Tab](#) | 769

Edit Container LSP Parameters

To edit the parameters of a container LSP:

1. Click **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. Hover over the **More Tabs** list and select **Container LSP**.

The Container LSP tab appears.

3. Select the container LSP you want to edit and click the edit (pencil) icon.

The Edit Container LSP page appears on the right.

4. Configure the container LSP as per [Fields on the Add Container LSP Page on page 770](#).

NOTE: You cannot edit the following fields:

On the **Properties** tab,

- Provisioning Method
- Provision Type
- Container Name
- Node A
- Node Z
- On the **Advanced** tab, IP Z.

5. Click **Edit**.

A confirmation message appears stating that the Edit Container LSP request is successfully sent.

RELATED DOCUMENTATION

| [Add a Container LSP | 770](#)

Maintenance Event Overview

You can schedule maintenance events for network elements (nodes, links, or facilities), so that you can perform updates or other configuration tasks.

A maintenance event is a planned downtime which occurs at a scheduled future date and time. During a scheduled maintenance event, the selected elements are considered logically down, and Paragon Automation reroutes the LSPs around those elements.

NOTE: Paragon Automation only attempts to reoptimize PCE-initiated and PCC-delegated LSPs (not PCC-controlled LSPs). PCC-controlled LSPs are not rerouted to avoid scheduled maintenance events.

The maintenance event that you add is listed under the **Maintenance** tab on the **Topology** page. You can view the progress of the maintenance event from its status. A maintenance event can have one of the following status:

- **Planned**—Event is scheduled some time in the future.
- **Completed**—Event is completed.
- **In Progress**—Event is in progress.
- **Canceled**—The scheduled event has been canceled.

After the maintenance event is completed, by default, all LSPs that were affected by the event are optimized again and several maintenance reports are generated that can be viewed from **Reports > Maintenance > Report-Name** page.

You can view, add, simulate, edit, cancel, or delete the scheduled maintenance events for network elements from the **Maintenance** tab of the network information table (**Network > Topology**). Hover over the **More** Tabs list and select **Maintenance**. When a network element (node, link, or facility) is undergoing a maintenance event, it appears on the topology map with a red M (for maintenance).

NOTE: You cannot delete a maintenance event that is in progress. You can, however, cancel one. You might want to cancel an event rather than delete it if you think you will reactivate it later, possibly with modifications.

RELATED DOCUMENTATION

[Understand How Pathfinder Handles LSPs | 715](#)

[Add a Maintenance Event | 778](#)

[Maintenance Reports Overview | 918](#)

About the Maintenance Tab

IN THIS SECTION

- [Tasks You Can Perform | 777](#)

You can use the Maintenance tab in the network information table to schedule maintenance events for network elements, so that you can perform updates or other configuration tasks.

NOTE: Paragon Pathfinder only attempts to reoptimize PCE-initiated and PCC-delegated LSPs (not PCC-controlled LSPs).

Tasks You Can Perform

You can perform the following tasks:

- Add a maintenance event. See ["Add a Maintenance Event" on page 778](#).
- Edit an existing maintenance event. See ["Edit a Maintenance Event" on page 780](#).
- Delete a maintenance event. See ["Delete a Maintenance Event" on page 782](#).
- Simulate a maintenance event. See ["Simulate a Maintenance Event" on page 781](#).
- View health of a device in maintenance—Select the maintenance event and click **Device Health**. Select the device for which you want to view the health information. You are redirected to the Network Health page, which displays the device health information in a table and tile view. For more information, see ["About the Network Health Page" on page 833](#).
- Download maintenance event information—Click **Download** to download information about all planned maintenance events in CSV format.
- View Details—View detailed information about the existing maintenance events by clicking **More >Show Detail** or click the Details icon displayed next to the maintenance event name when you hover over it.

RELATED DOCUMENTATION

[About the Topology Page | 677](#)

[Interactive Map Features Overview | 660](#)

Add a Maintenance Event

To add a maintenance event:

1. Click **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. Hover over the **More Tabs** list and select **Maintenance**.

The Maintenance tab appears.

3. Click the add (+) icon.

The Add Maintenance page appears on the right.

4. Configure the maintenance event as per [Table 118 on page 778](#).

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

5. Click **Add**.

A confirmation message appears stating that the add maintenance request is sent successfully. The new maintenance event is displayed under the **Maintenance** tab.

NOTE: When an element (node, link, or SRLG) is undergoing maintenance, it appears on the topology map with a red M (for maintenance).

Table 118: Fields on the Add Maintenance Page

Field	Description
Name	Enter a name for the maintenance event. NOTE: Ensure that there are no spaces in the maintenance event name.

Table 118: Fields on the Add Maintenance Page (Continued)

Field	Description
Start	<ol style="list-style-type: none"> 1. Click the calendar icon on the right to display a calendar from which you can select the year, month, and day. 2. Click Select Time to enter a custom start time by scrolling up or down to select the hour, minute, and second along with AM/PM selection. NOTE: You can also click Now to select the current date and time. 3. Click Ok. NOTE: You can manually enter the date and time values.
End	<p>Enter an end time for the maintenance event (similar to Start time).</p> <p>NOTE: The minimum time interval between start and end time of a maintenance event should be 5 minutes.</p>
Owner	Auto-populates with the name of the user who is scheduling the maintenance event.
Description	Enter a description for the maintenance event.
Auto Complete at End Time	<p>Automatically completes the maintenance event and brings the elements back up at the specified end time. If the check box is not selected, you must manually complete the maintenance event after it finishes.</p> <p>When a maintenance event is completed, it brings the maintenance elements back to an Up state, ready for path reoptimization. The affected LSPs are then rerouted to optimal paths unless you have deselected Reoptimize Tunnels Upon Completion.</p>
Reoptimize Tunnels upon Completion	<p>When enabled (the default behavior), Paragon Pathfinder calculates and moves the LSP to the most optimum path after the maintenance event is complete.</p> <p>When disabled, Paragon Pathfinder does not calculate the most optimum path and lets the LSP remain in its current path after the maintenance event is complete.</p>
Nodes	Select the nodes in the left column and click the right arrow to move them to the right (selected) column. Click the left arrow to deselect elements. You can also search for a specific node from Search here option.

Table 118: Fields on the Add Maintenance Page (Continued)

Field	Description
Links	Select the links in the left column and click the right arrow to move them to the right (selected) column. Click the left arrow to deselect elements. You can also search for a specific link from Search here option.
Facilities	Select the facilities in the left column and click the right arrow to move them to the right (selected) column. Click the left arrow to deselect elements. You can also search for a specific facility from Search here option.

RELATED DOCUMENTATION

| [About the Maintenance Tab | 777](#)

Edit a Maintenance Event

To edit a maintenance event:

1. Click **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. Hover over the **More Tabs** list and select **Maintenance**.

The Maintenance tab appears.

3. Select the maintenance event you want to edit.

4. Click the edit (pencil) icon.

The Edit Maintenance page appears on the right.

NOTE: You cannot edit the parameters of a maintenance event which is in progress. You can only edit the **Status** of the event.

5. Configure the maintenance event as per [Table 119 on page 781](#).

6. Click **Edit**.

A confirmation message is displayed stating that the edit maintenance request is sent successfully.

Table 119: Fields on the Edit Maintenance Page

Field	Description
Status	<p>Edit the status of the maintenance event by selecting one of the following:</p> <ul style="list-style-type: none"> • Planned—The maintenance event execution is planned. • Completed—The maintenance event execution is completed. • Cancelled—When you cancel a maintenance event, it remains in the Maintenance tab of the network information table, with the operation status as Cancelled. You might want to cancel an event rather than deleting it, as you can reactivate it later with possible modifications, or use it for tracking purposes. • Deleted—When you delete an event, it is removed from the network information table. <p>NOTE: You cannot delete a maintenance event that is in progress. You can, however, cancel one.</p> <p>When an element (node, link, or SRLG) is undergoing a maintenance event, it appears on the topology map with a red M (for maintenance).</p>

NOTE: For information on other parameters, see [Fields on the Add Maintenance Page on page 778](#).

RELATED DOCUMENTATION

[About the Maintenance Tab | 777](#)

Simulate a Maintenance Event

You can run scheduled maintenance event simulations for different failure scenarios to test the resilience of your network. Network simulation is based on the current network state for the selected maintenance events at the time the simulation is initiated. This simulation workflow can not simulate a maintenance event for a future network state or simulate elements from other concurrent maintenance events. You can run network simulations based on elements selected for a maintenance event, with the option to include exhaustive failure testing.

To simulate a maintenance event:

1. Click **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. Hover over the **More** Tabs list and select **Maintenance**.

The Maintenance tab appears.

3. Select the maintenance event you want to simulate.

4. Select **Simulate**.

The Simulate Maintenance page appears. Exhaustive Failure simulation options are displayed along with the nodes, links, and facilities associated with the maintenance event.

5. (Optional)—Select **Nodes**, **Links**, or **Facilities** check box if you want to run an Exhaustive Failure Simulation on those network elements.

If you do not perform an exhaustive failure simulation (all check boxes under Exhaustive Failure Simulation are unchecked), all the nodes, links, and facilities selected for the maintenance event fail concurrently (at the same time). For example, if Nodes are selected under the Exhaustive Failure Simulation option, the simulation will still fail all the maintenance event elements concurrently, but fails each of the other nodes in the topology, one at a time. If you select multiple element types for exhaustive failure simulation, all possible combinations involving those elements are tested. The subsequent report reflects peak values based on the worst performing combination.

6. Click **Simulate** to perform the failure simulation and generate reports.

A confirmation message appears stating that the simulation is complete and the generated reports are available in the Reports (**Reports > Maintenance**) section.

RELATED DOCUMENTATION

| [About the Maintenance Tab | 777](#)

Delete a Maintenance Event

To delete a maintenance event:

1. Click **Network > Topology**.

The topology map is displayed with network information table at the bottom of the page.

2. On the **Maintenance** tab, select the maintenance event you want to delete.

3. Click the Delete (trash can) icon.

A confirmation message appears.

4. Click **Yes**.

The maintenance event is deleted and removed from the network information table. The topology map is updated.

NOTE: You cannot delete a maintenance event when its **Status** is **In progress**. You must edit the maintenance event status by canceling or completing the event before you can delete it. For more information on event status, see [Fields on the Edit Maintenance Page on page 781](#).

About the P2MP Groups Tab

IN THIS SECTION

- [PCEP-Provisioned P2MP Groups with Service Mapping | 783](#)
- [Required Router Configuration | 784](#)
- [Automatic Rerouting Around Points of Failure | 786](#)
- [P2MP Tree Design with Diverse PE to CE Links | 786](#)
- [Tasks You Can Perform | 790](#)

You can use the **P2MP Groups** tab in the network information table on the Topology (**Network > Topology**) page to view, configure, and manage the point-to-multipoint group (P2MP) in your network.

P2MP groups, or trees, can be provisioned to help conserve bandwidth. Bandwidth is replicated at branch points. P2MP groups support NETCONF and PCEP provisioning methods. PCEP provisioning offers an advantage of real-time reporting. For PCEP-provisioned P2MP groups, Paragon Pathfinder also supports association with a multicast VPN instance (S,G) flow in its PCEP P2MP service mapping functionality. It's important to understand the effect on the network flows when you make various configuration changes on the router, so we recommend using your Junos OS documentation as a reference.

PCEP-Provisioned P2MP Groups with Service Mapping

Beginning with Junos OS Release 19.4R1, Junos OS has the ability to associate multicast flows (S,G) in the multicast VPN context to a PCEP P2MP LSP provisioned through the PCE, in accordance with *draft-*

ietf-pce-pcep-flowspec-05. You can leverage that Junos OS functionality by provisioning PCEP P2MP groups that you associate with one or more multicast flows (S,G) in a multicast VPN. Once a P2MP group is associated with a particular (S,G) in a multicast VPN, traffic from that particular source IP S going to group IP G, is able to utilize that P2MP group.

The PCE side of this functionality requires specific configuration on the router which is described in the following section.

NOTE: This service mapping functionality is not the same as the user properties-based service mapping that is supported for NETCONF-provisioned LSPs.

Required Router Configuration

In Junos OS Release 15.1F6 and later, you can configure the router to automatically send P2MP LSP information to a controller in real time. Without that configuration, you must run device collection for PCE to learn about newly provisioned P2MP LSPs. The configuration is done in the [set protocols pcep] hierarchy for PCEs and PCE groups.

The following configuration statement allows PCEP to report the status of P2MP trees in real time:

```
set protocols pcep pce pce-id p2mp-lsp-report-capability
```

For PCEP-provisioning, the following additional configuration statements are also required:

```
set protocols pcep pce pce-id p2mp-lsp-update-capability
set protocols pcep pce pce-id p2mp-lsp-init-capability
```

For PCEP P2MP service mapping with flowspec, the following additional configuration statements are required:

- ```
set protocols pcep pce pce-id pce-traffic-steering
```

This configuration enables the router to support traffic steering (flowspec), and must be configured on the head-end.

- ```
set routing-instances routing-instance-name provider-tunnel external-controller pccd
set routing-instances routing-instance-name protocols mvpn sender-site
```

This configuration enables the router to accept P2MP LSP provisioning with (S,G) for this multicast VPN from an external controller.

Sample configuration on a router with a P2MP group head-end is as follows:

```

set protocols pcep pce pce-id p2mp-lsp-report-capability
set protocols pcep pce pce-id p2mp-lsp-update-capability
set protocols pcep pce pce-id p2mp-lsp-init-capability
set protocols pcep pce pce-id pce-traffic-steering
set routing-instances routing-instance-name routing-options static route ip-address next-hop ip-address
set routing-instances routing-instance-name routing-options multicast ssm-groups multicast-address-group
set routing-instances routing-instance-name protocols pim interface interface mode sparse
set routing-instances routing-instance-name protocols mvpn sender-site
set routing-instances routing-instance-name instance-type vrf
set routing-instances routing-instance-name provider-tunnel external-controller pccd
set routing-instances routing-instance-name provider-tunnel rsvp-te label-switched-path-template mvpn-template
set routing-instances routing-instance-name interface interface
set routing-instances routing-instance-name route-distinguisher RD
set routing-instances routing-instance-name vrf-target route-target
set routing-instances routing-instance-name vrf-table-label
set protocols mpls label-switched-path mvpn-template template p2mp

```

NOTE: After provisioning P2MP LSPs, there can be a PCEP flap which can cause the UI display for RSVP utilization and RSVP live utilization to be out of sync. This is also true for P2P LSPs. You can display utilization metrics by navigating to **Performance** in the left widget of the UI. This is a UI display issue only. The next live update from the network or the next manual sync using **Sync Network Model (Configuration > Network > Pathfinder)** corrects the UI display.

Useful Junos OS show commands

[Table 120 on page 786](#) describes the Junos OS commands to view detailed information about P2MP groups and their associated sub-LSPs.

Table 120: Junos OS show commands for P2MP Groups

Junos OS Command	Description
<code>show mpls lsp p2mp ingress extensive</code>	Displays information on each sub-LSP that belongs to the P2MP group.
<code>show path-computation-client lsp</code>	Displays information about LSPs, including their status and type.
<code>show rsvp session p2mp ingress statistics</code>	Displays whether the selective tunnel is taking precedence over the original inclusive dynamic provider tunnel as expected.
<code>show mvpn c-multicast instance-name <i>mvpn-instance-name</i> display-tunnel-name</code>	Displays which P2MP group is mapped to which (S,G) within a multicast VPN.
<code>show path-computation-client traffic-steering</code>	Displays all P2MP groups with flowspec ID, state, and (S,G) prefix.

Automatic Rerouting Around Points of Failure


For PCEP-provisioned P2MP groups only (including those with flow-mapping information), sub-LSPs are dynamically rerouted around points of failure along the path of the tree. You might not see the Op Status in the network information table change during the reroute as it happens very quickly. The topology map displays a red F on any failed link or node, and you can see how the path is rerouted around those markers.

P2MP Tree Design with Diverse PE to CE Links

Paragon Pathfinder can calculate diverse P2MP tree designs all the way to a customer edge (CE) node or site. Although the path computation extends to the intended endpoints, CE nodes or sites, the sub-LSPs actually terminate on the provider edge (PE) nodes. This ensures that the selection of tail nodes best satisfies the diversity constraints. When the diverse P2MP trees are computed, Paragon Pathfinder considers the shared risk groups and affinity constraints in the PE-CE links.

[Figure 60 on page 787](#) shows an example topology with P2MP tree design. Two data sources are located at PE nodes (PE1 and PE2). The receiving nodes are located in sites A, B, C, and D with each site having two CE nodes.


Figure 60: P2MP Tree Design with Diverse PE to CE Links Example Topology



Calculating pairs of diverse P2MP LSPs that terminate in the same nodes or sites enables the two data sources to distribute redundant data streams to the CE nodes or sites.


Sites A, B, C, and D could also be depicted as shown in [Figure 61 on page 788](#), where each site is shown as a pseudo-node. In this context, a pseudo-node is any node or site not discovered or managed by Paragon Pathfinder. The pseudo-nodes can be used as the leaf nodes in tree calculations.

Figure 61: Diverse PE to CE Links Topology with Pseudo-Nodes



The goal is for each sub-LSP in a diverse pair to be routed over a different PE-CE link, transiting a different PE node. Each sub-LSP terminates on the PE node that is the penultimate hop to the destination. An example is shown in [Figure 62 on page 788](#). The redundant data streams are represented by the two colored paths.

Figure 62: Diverse PE to CE Links Topology with Redundant Data Streams



You can request a tree design with diverse PE to CE links through the REST API. When the tree is created, the Explicit Route Object (ERO) of the route, which is the list of LSRs specified by IP addresses through which the path must pass, is always strict. This is in contrast to a single P2MP group in which a loose path to the address can be specified. The result is that if a link goes down in a diverse P2MP tree, the signaling address remains the same and the operational status of the tunnel remains down. If the link comes back up, the tunnel is restored. You can also delete the tunnel or schedule a new one.


The Path Computation Server (PCS) follows a set of rules when routing LSPs to CE nodes or sites:

- The links connecting to the CE nodes or sites might not have protocols enabled since they do not need to support LSPs. But the PCS must still consider those links for path placement.
- The PCS will never use a pseudo-node marked as an Access node as a transit node. Access nodes are end destination nodes.
- LSPs to CE nodes or sites must be signaled to their penultimate hop, the final PE in the path.
- The ERO pushed to the devices must be trimmed so it does not include the pseudo-node addresses or the PE-CE link identifiers.

NOTE: Paragon Pathfinder cannot discover CE nodes and sites, so you must add them and their associated links to the topology, using either the UI or the REST API. For more information on adding nodes using the Paragon Automation UI, see ["Add a Node" on page 696](#) and ["Add a Link" on page 705](#).

[Figure 63 on page 790](#) shows a sample topology with some CE nodes that are configured as **Regular** and others that are configured as **Access**. For example, nodes CE_A1 and CE_B1 are Regular nodes because they need to be used for transit to the alpha and beta sites respectively. CE_D1 and CE_D2 are examples of Access nodes that represent end points, and are not used for transit purposes. In addition to creating the nodes, you must also add links between the CE nodes/sites and the PE nodes in the topology. As the nodes are pseudo-nodes, the links connecting them are pseudo-links. As pseudo-links, their status will remain Unknown in the **Status** column of the **Link** tab in the network information table.

Figure 63: Sample Topology with Access and Regular Pseudo-Nodes

**NOTE:**

- Hybrid diverse P2MP groups (destination designation includes both PEs and CEs) is not supported.
- You cannot provision diverse P2MP trees using the Paragon Pathfinder UI. It must be done through the REST API.
- Diverse P2MP with flow mapping is not supported. The workaround for this limitation is to create a diverse tree using the REST API and then assign mapping to individual P2MP groups using either the UI or the REST API.

Tasks You Can Perform

You can perform the following tasks:

- Add a P2MP Group. See ["Add a P2MP Group"](#) on page 791.

- Edit a P2MP Group. See ["Edit P2MP Group Parameters" on page 798](#).
- Delete a P2MP Group—On the **P2MP Group** tab of the network information table on the **Topology** page, select the P2MP group you want to delete and click the delete (trash can) icon. A confirmation message appears. Click **Yes**. The P2MP group is deleted from the network information table.

Alternatively, you can use the **Tunnel** tab of the network information table to delete all the sub-LSPs in the P2MP group, which also deletes the group itself.

NOTE: When you delete a P2MP group, all sub-LSPs that are part of that group are also deleted. If you delete a P2MP group with associated multicast flows (S,G) in a multicast VPN context, the flows are also deleted.

- View Sub-LSPs— From the **More** list, you can view Sub-LSPs associated with the P2MP group. When you click the **Sub-LSPs** option, you are redirected to the **Tunnel** tab where detailed information about the sub-LSPs is displayed. You can perform several tasks such as add, edit, or delete LSPs from the Tunnel tab. For more information, see ["About the Tunnel Tab" on page 710](#).
- Download P2MP Groups information—Click **Download** to export (download) information about all the P2MP groups in CSV format.
- View detailed information about the P2MP group by clicking **More > Show Detail** or click the Details icon which appears next to the P2MP group name when you hover over it. A pop-up appears displaying the P2MP tree, flows associated with the group, and link/LSP details that are a part of the selected P2MP group.

RELATED DOCUMENTATION

[About the Topology Page | 677](#)

[Interactive Map Features Overview | 660](#)

Add a P2MP Group

To add a P2MP Group:

1. Click **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. Hover over the **More** Tabs list and select **P2MP Group**.

The P2MP Group tab appears.

3. Click the add (+) icon.

The Add P2MP Group page appears on the right.

4. Configure the P2MP group as per [Table 121 on page 792](#).

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

5. Click **Add**.

A confirmation message appears stating that the Add P2MP request is successfully sent. The new P2MP group is displayed under the **P2MP Groups** tab.

Table 121: Fields on the Add P2MP Group Page

Field	Description
Properties	
Provisioning Method	Select PCEP or NETCONF. The default is NETCONF.
Provision Type	The default is RSVP, which is the only option supported for P2MP groups.
P2MP Group Name	Enter a user-defined name for the P2MP group. Only alphanumeric characters, hyphens, and underscores are allowed. Other special characters and spaces are not allowed.
Node A	Click the Node A field and select the name of the ingress (source) node from the list.
Node Z List	Click the Node Z field and select one or more egress (destination) nodes from the list.
Planned Bandwidth	Enter the planned bandwidth value. You must enter a number immediately followed by K (Kbps), M (Mbps), or G (Gbps). For example, 10M signifying 10 megabits per second. NOTE: If you enter a value without units, bps is applied.
Planned Metric	Enter the metric value for static tunnel. Use the up and down arrows to increment or decrement the value. For more information, see Basic LSP Configuration in <i>MPLS Applications User Guide</i> .

Table 121: Fields on the Add P2MP Group Page (Continued)

Field	Description
Setup	Enter the RSVP 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 LSP definition in Junos OS.
Hold	Enter the RSVP 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 LSP definition in Junos OS.
Routing Method	<p>Click the Routing Method list and select a routing method from the following:</p> <ul style="list-style-type: none"> • Admin Weight • Constant • Default <p>NOTE: Do not change the routing method for PCEP-provisioned sub-LSPs; they should always have a routing method of “default”.</p> <ul style="list-style-type: none"> • Delay • Distance • ISIS • OSPF • Route by Device <p>NOTE: For NETCONF-provisioned P2MP group, the default routing method is Route By Device and the path for all sub-LSPs is dynamic.</p> <p>For PCEP-provisioned P2MP group, default is selected as the routing method. The Route By Device routing method is not available for PCEP-provisioned P2MP group.</p>

Table 121: Fields on the Add P2MP Group Page (Continued)

Field	Description
Bandwidth Sizing	<p>Toggle the button to enable bandwidth sizing. By default, this option is disabled.</p> <p>NOTE: You cannot enable Bandwidth Sizing if the provisioning method is NETCONF.</p> <p>Once enabled, configure the following parameters:</p> <ul style="list-style-type: none"> • Adjustment threshold (in percentage): Controls the sensitivity of the automatic bandwidth adjustment. The new planned bandwidth is only considered if it differs from the existing bandwidth by the value of this setting or more. • Minimum and Maximum (planned) Bandwidth: <ul style="list-style-type: none"> If the new planned bandwidth is greater than the maximum setting, PathFinder signals the LSP with the maximum bandwidth. If the new planned bandwidth is less than the minimum setting, PathFinder signals the LSP with the minimum bandwidth. If the new planned bandwidth value is in between the maximum and minimum settings, Pathfinder signals the LSP with the new planned bandwidth. • Minimum Variation Threshold: Specifies the sensitivity of the automatic bandwidth adjustment when the new planned bandwidth is compared to the current planned bandwidth. The new planned bandwidth is only considered if the difference is greater than or equal to the value of this setting. This can be used to prevent small fluctuations from triggering unnecessary bandwidth changes. If both the adjustment threshold and the minimum variation threshold are greater than zero, both settings are taken into consideration. In that case, the new planned bandwidth is considered if: <ul style="list-style-type: none"> • The percentage difference is greater than or equal to the adjustment threshold and • The actual difference is greater than or equal to the minimum variation.
Constraints	
Enables you to configure administrative groups (link coloring) for bypass LSPs. For more information, see "Assign Names to Admin Group Bits" on page 197 .	
Admin Group Include All	Specifies the administrative groups whose links the bypass LSP must traverse. Select one or more bit-level link coloring options from the list.

Table 121: Fields on the Add P2MP Group Page (Continued)

Field	Description
Admin Group Include Any	Specifies the administrative groups whose links the bypass LSP can traverse. Select one or more bit-level link coloring options from the list.
Admin Group Exclude	Specifies the administrative groups to exclude for a bypass LSP. Select one or more bit-level link coloring options from the list.
Advanced	
Description	Enter a description for the P2MP Group.
Diversity Group	Enter the name of a group of tunnels to which this tunnel belongs, and for which diverse paths is required.
Diversity Level	<p>Select one of the following values for diversity level:</p> <ul style="list-style-type: none"> • Default (no diversity) • Link • Shared Risk Link Group (SRLG) • Site <p>Site diversity is the strongest—it includes SRLG and link diversity. SRLG diversity includes link diversity. Link diversity is the weakest.</p>

Table 121: Fields on the Add P2MP Group Page (Continued)

Field	Description
Custom Attributes	<p>Simple name-value pairs which can be used to add any arbitrary customer-specific information. For example, to differentiate between properties for different vendor nodes.</p> <p>To add custom properties associated with the node:</p> <ol style="list-style-type: none"> 1. Click add (+) icon to add a new row. 2. Click the newly added row to enter the Name and Value. 3. Click the ✓ icon to save your changes. <p>NOTE:</p> <ul style="list-style-type: none"> • You can add multiple rows (properties). • To delete an entry (row), select the row and click the delete (trash can) icon.
<p>Service</p> <p>NOTE: This tab is available only when the provisioning method selected is PCEP. When provisioning method is NETCONF, this tab is unavailable (greyed out).</p>	
VPN Identifier	<p>Select one of the following:</p> <ul style="list-style-type: none"> • MVPN instance: Select the multicast VPN instance as configured on the router. • Route Distinguisher (RD): Populates automatically when you select the multicast VPN instance, if the information is available. If the field does not automatically populate, you can: <ul style="list-style-type: none"> • Obtain the RD from the head-end router configuration. • Run device collection for the head-end router from the Node tab of the network information table. Right-click the head-end router and select More > Run Device Collection. Once the collection completes, the RD field should auto-populate.

Table 121: Fields on the Add P2MP Group Page (Continued)

Field	Description
Flows	<p>Define the (S,G) groups.</p> <ol style="list-style-type: none"> 1. For each flow you want to add, click the add (+) icon to display a new line. 2. Enter the source and group (multicast traffic destination) IP addresses and mask. <p>NOTE: To delete an entry, select the row and click the delete (trash can) icon.</p> <p>The maximum number of flows (S,G groups) that you can successfully associate with a P2MP LSP depends on the traffic rate and group type (SSM or ASM), and is therefore, not a fixed number. If you surpass the maximum number of flows for the P2MP group, all the flows go into "Inactive (mapping successful)" state as viewed in the output of the <code>show path-computation-client traffic-steering</code> command on the router. To recover, delete the P2MP group and recreate it.</p>
Schedule	
Plan	<p>Select one of the following options:</p> <ul style="list-style-type: none"> • No Schedule • Once: Enter start and end date and time for a single event. <ol style="list-style-type: none"> 1. Click the calendar icon on the right to display a monthly calendar from which you can select the year, month, and day. 2. Click Select Time to enter a custom start time by scrolling up or down to select the hour, minute, and second along with AM/PM selection. <p>NOTE: You can also click Now to select the current date and time.</p> <ol style="list-style-type: none"> 3. Click Ok. <p>NOTE: You can also manually enter date and time values.</p> • Recurring Daily: Enter the start and end parameters for a recurring daily event (similar to the Once schedule).

RELATED DOCUMENTATION

[About the P2MP Groups Tab](#) | 783

Edit P2MP Group Parameters

To edit the parameters of a P2MP group:

1. Click **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. Hover over the **More** Tabs list and select **P2MP Group**.

The P2MP Group tab appears.

3. Select the P2MP group that you want to edit and click the edit (pencil) icon.

The Edit P2MP Group page appears on the right.

4. Configure the P2MP group parameters as per [Fields on the Add P2MP Group Page on page 792](#).

NOTE:

- You cannot edit the Provisioning Method, Provision Type, P2MP Group Name, and Node A field values.
- The Service tab is only displayed if you are modifying a PCEP-provisioned P2MP group with service mapping.

5. Click **Edit**.

A confirmation message appears stating that the edit P2MP Group request is successfully sent.

NOTE: If the sub-LSPs tab in the network information table fails to update after modifying P2MP group or sub-LSP attributes, you can close the sub-LSPs tab and reopen it to refresh the display. You can also use the refresh button at the bottom of the table that turns orange when prompting you for a refresh. When you click the refresh button, the Web UI client retrieves the latest P2MP sub-LSP status from the server.

RELATED DOCUMENTATION

[About the P2MP Groups Tab](#) | 783

About the SRLG/Facility Tab

IN THIS SECTION

- [Tasks You Can Perform | 799](#)

Shared Link Risk Group (SRLG) information can be received from BGP-LS or Transport controller, whenever a path optimization occurs or whenever some event triggers rerouting. The information from these sources is merged and displayed on the UI. You can add, edit, and delete user-defined SRLGs from the **SRLG/Facility** tab of the network information table on the Topology (**Network**> **Topology**) page.

Transport SRLG information is considered whenever a path optimization occurs or whenever some event triggers rerouting. By default, Paragon Automation associates transport SRLGs to IP links based on information received from the transport controller. Connecting to more than one transport controller introduces the possibility of overlapping SRLG ranges, which might not be acceptable. The configuration of transport controller profiles allows for the specification of an additional `TSRLGprefix` (a prefix extension) for each transport controller to prevent unintentional overlap. Preventing unintentional SRLG range overlap requires vigilance when you have transport controller ranges and you also manually assign SRLGs to IP links.

Tasks You Can Perform

You can perform the following tasks:

- Add an SRLG/Facility. See ["Add an SRLG/Facility" on page 800](#).
- Edit existing SRLG/Facility parameters. See ["Edit SRLG/Facility Parameters" on page 801](#).
- Delete an existing SRLG/Facility—On the **SRLG/Facility** tab of the network information table in the **Topology** page, select the facility that you want to delete and click the delete (trash can) icon. A confirmation message appears. Click **Yes**. The facility is deleted from the network information table.
- View details about the SRLG/Facility by clicking the details icon when you hover over the SRLG/facility name or click **More** > **Show Detail**.

RELATED DOCUMENTATION

[About the Topology Page | 677](#)

Add an SRLG/Facility

To add an SRLG/Facility:

1. Click **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. Hover over the **More Tabs list and select **SRLG/Facility**.**

The SRLG/Facility tab appears.

3. Click the add (+) icon.

The Add Facility page appears.

4. Configure the fields as per [Table 122 on page 800](#).

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

5. Click **Add.**

A confirmation message appears stating that the add facility request has been successfully sent.

Table 122: Fields on the Add Facility Page

Field	Description
Name	Enter a name for the SRLG/Facility.
Nodes	Select the nodes in the left column and click the right arrow to move them to the right (selected) column. Click the left arrow to deselect elements. You can also search for a specific node from Search here field.
Links	Select the links in the left column and click the right arrow to move them to the right (selected) column. Click the left arrow to deselect elements. You can also search for a specific link from Search here field.

RELATED DOCUMENTATION

| [About the SRLG/Facility Tab | 799](#)

Edit SRLG/Facility Parameters

To edit the parameters of an SRLG/Facility:

1. Click **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. Hover over the **More** Tabs list and select **SRLG/Facility**.

The SRLG/Facility tab appears.

3. Select the facility you want to edit and click the edit (pencil) icon.

The Edit Facility page appears.

4. Configure the fields as per [Fields on the Add Facility Page on page 800](#).

5. Click **Edit**.

A confirmation message appears stating that the edit facility request is successfully sent.

RELATED DOCUMENTATION

| [About the SRLG/Facility Tab | 799](#)

About the Topology Group Tab

IN THIS SECTION

- [Tasks You Can Perform | 802](#)

A topology group, also referred to as a topology slice, is a subset of the nodes and links in a network. A topology group is identified by a slice ID. You must configure the **Slices** parameter for the required nodes or links to group them as a topology group.

You can configure the nodes and links in a network to form a topology group in one of the following ways:

- Edit the node and link parameters to specify one or more slice IDs (positive integers) in the **Slices** field. For more information, see ["Edit Node Parameters" on page 699](#) and ["Edit Link Parameters" on page 708](#).
- Select the required nodes and links on the topology map to add a topology group. For more information, see ["Group Nodes and Links into a Topology Group" on page 683](#).

All the network elements that are tagged with the same slice ID form a topology group and are listed under the Topology Group tab. You can use topology groups as a constraint to restrict tunnel routing within a topology slice. For more information, see ["Add a Single LSP" on page 729](#).

When you select a topology group in the network information table, the nodes and links within the topology group are highlighted in yellow on the topology map.

Tasks You Can Perform

You can perform the following tasks on the Topology Group tab:

- Hide unrelated nodes and links.

Select **Hide Unrelated** to show only the nodes and links within the selected topology group on the topology map.

- Edit the description for the topology group.

Select the topology group and click **Edit** (pencil icon). The Edit topology Group page appears. In the **Description** field, enter a description for the topology group, and click **Edit**. The updated description of the topology group is displayed on the network information table.

- Delete a topology group.

Select the topology group, and click the **Delete** (trash can) icon. On the confirmation message that appears, click **Yes**. The topology group is removed from the network information table.

NOTE: Before you delete a topology group, you must remove all the nodes and links within the topology group. To remove a node or link from a topology group, you must remove the slice ID value from the **Slices** parameter for the particular node or link. For more information, see ["Edit Link Parameters" on page 708](#) and ["Edit Node Parameters" on page 699](#).

- View topology group details.

From the **More** list, select **Show Detail** to view details about the nodes and links that are part of the topology group.

RELATED DOCUMENTATION

[About the Topology Page | 677](#)

[Interactive Map Features Overview | 660](#)

Add Anycast Group Tunnels

You can add a segment routing tunnel with an anycast group as the destination, so that the traffic is routed to any one of the nodes that is part of the anycast group.

To add a segment routing tunnel with an anycast group as the destination:

1. Select [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. In the Anycast Group tab, select [Provisioning > Add Anycast Group Tunnel](#).

The Add Anycast Group Tunnel page appears.

3. Complete the configuration on each tab according to the guidelines in [Table 123 on page 804](#).

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

4. (Optional) You can click [Preview Path](#) at the bottom of the page to see the path drawn on the topology map.

5. Click [Add](#) to add the tunnel.

A confirmation message appears on the top of the page, indicating that a provision anycast tunnel request was successfully created:

- If you have the Auto-approve permission assigned to your user role, the request is automatically approved and deployed on the devices.
- If you don't have Auto-approve permission, the request must be manually approved and then, deployed by a user having the required permissions. See "[About the Change Control Management Page](#)" on page 829.

The tunnel then appears in Tunnel tab of the network information table (in the Topology page).

Table 123: Fields on the Add Anycast Group Tunnel Page

Field	Description
<i>Properties</i>	
Provisioning Method	<p>From the list, select one of the following methods to be used to provision the tunnel:</p> <ul style="list-style-type: none"> • NETCONF (default)—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. • PCEP (Path Computation Element Protocol)—The tunnel is initiated by the path computation element (PCE) and the associated configuration statements do not appear in the router configuration file. Upon provisioning, this tunnel is added as a PCE-initiated tunnel. <p>NOTE:</p> <ul style="list-style-type: none"> • For IOS-XR routers, NETCONF-based tunnel provisioning has the same capabilities as PCEP-based tunnel provisioning.
Provision Type	Displays the type of tunnel to be provisioned—SR (segment routing).
Name	<p>Specify a unique name for the tunnel.</p> <p>You can use any number of alphanumeric characters, hyphens, and underscores.</p>
Node A	From the list, select the node that you want to use as the ingress node.
Node Z	<p>Do one of the following:</p> <ul style="list-style-type: none"> • If you want the traffic to be routed to a single node, select that node as the egress node. • If you want the traffic to be routed to one of the nodes in an anycast group, click the toggle button (Select destination as anycast group) and select the anycast group from the list. By default, this toggle button is disabled.

Table 123: Fields on the Add Anycast Group Tunnel Page (Continued)

Field	Description
Admin Status	<p>The Path Computation Server (PCS) uses the administration status of the tunnel to decide whether to route or provision, or both route and provision the tunnel.</p> <p>If the tunnel is routed, no traffic flows through the tunnel and its operational status is Up. If the tunnel is provisioned, traffic flows through the tunnel and its operational status is Active.</p> <p>Select one of the following options as the administration status:</p> <ul style="list-style-type: none"> • Up—If you select this option, the PCS routes and provisions the tunnel. • Planned—If you select this option, the PCS routes the tunnel and reserves capacities for the tunnel. However, the PCS doesn't provision the tunnel. • Shutdown—If you select this option, the PCS neither routes nor provisions the tunnel. The tunnel is maintained in the datastore and is associated with a persist state so that the tunnel can be brought back up at a later time, if required.
Path Type	From the list, select primary, secondary, or standby as the path type.
Path Name	Specify the name for the path.
Planned Bandwidth	<p>Specify the planned bandwidth (along with valid units, with no space between the bandwidth and units) for the tunnel.</p> <p>If you specify a value without units, bps is applied.</p> <p>Valid units are:</p> <ul style="list-style-type: none"> • B or b • M or m • K or k • G or g <p>Examples: 50M, 1000b, 25g.</p>

Table 123: Fields on the Add Anycast Group Tunnel Page (Continued)

Field	Description
Bandwidth Sizing	<p>NOTE: This option is displayed only when you select PCEP as the provisioning method.</p> <p>Click the toggle button to enable or disable (default) bandwidth sizing for the tunnel.</p> <p>If you enable bandwidth sizing, the tunnel is included in the periodic re-computation of planned bandwidth based on aggregated tunnel traffic statistics.</p>
Adjustment Threshold (%)	<p>NOTE: This option is available only when you enable bandwidth sizing.</p> <p>Specify the sensitivity (in %) of the automatic bandwidth adjustment.</p> <p>The new planned bandwidth is only considered if it differs from the existing bandwidth by the value of this setting or more. The default value is 10%.</p>
Minimum Bandwidth	<p>NOTE: This option is available only when you enable bandwidth sizing.</p> <p>Specify the minimum 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 applied.</p> <p>If the new planned bandwidth is less than the minimum setting, the PCE signals the tunnel with the minimum bandwidth. However, if the new planned bandwidth falls in between the maximum and minimum settings, The PCE signals the tunnel with the new planned bandwidth.</p> <p>The valid units are:</p> <ul style="list-style-type: none"> • B or b • M or m • K or k • G or g <p>Examples: 50M, 1000b, 25g.</p>

Table 123: Fields on the Add Anycast Group Tunnel Page (Continued)

Field	Description
Maximum Bandwidth	<p>NOTE: This option is available only when you enable bandwidth sizing.</p> <p>Specify the maximum 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 applied.</p> <p>If the new planned bandwidth is greater than the maximum setting, the PCE signals the tunnel with the maximum bandwidth. However, if the new planned bandwidth falls in between the maximum and minimum settings, the PCE signals the tunnel with the new planned bandwidth.</p> <p>The valid units are:</p> <ul style="list-style-type: none"> • B or b • M or m • K or k • G or g <p>Examples: 50M, 1000b, 25g.</p>
Minimum Variation Threshold	<p>NOTE: This option is available only when you enable bandwidth sizing.</p> <p>Specify the sensitivity of the automatic bandwidth adjustment when the new planned bandwidth is compared to the current planned bandwidth.</p> <p>Default: Zero.</p> <p>The new planned bandwidth is only considered if the difference is greater than or equal to the value of this setting. Because it is not a percentage, this can be used to prevent small fluctuations from triggering unnecessary bandwidth changes.</p> <p>If both the adjustment threshold and the minimum variation threshold are greater than zero, both settings are taken into consideration. In that case, the new planned bandwidth is considered if the percentage difference is greater than or equal to the adjustment threshold, and, the actual difference is greater than or equal to the minimum variation.</p>
Color Community	<p>NOTE: This field is available only for segment routing tunnels.</p> <p>Assign a color for the segment routing tunnel that can be used to map traffic on the tunnel.</p>

Table 123: Fields on the Add Anycast Group Tunnel Page (*Continued*)

Field	Description
Use Penultimate Hop as Signaling Address	<p>NOTE: This field is available only for segment routing tunnels.</p> <p>Click the toggle button to enable the PCS to use the penultimate hop as the signaling address for Egress Peer Engineering (EPE).</p> <p>If you haven't specified a color community, the setting applies to all traffic. If you've specified a color community, the setting applies to traffic in that color community.</p>
Setup	<p>Specify the setup priority for the tunnel traffic.</p> <p>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.</p> <p>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.</p>
Hold	<p>Specify the hold priority for the tunnel traffic.</p> <p>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.</p> <p>Based on the hold priority, the PCE determines whether the tunnel can be preempted or not. If the hold priority for an tunnel is higher, it is unlikely for the tunnel to be preempted.</p>
Planned Metric	<p>Specify the static tunnel metric.</p> <p>The PCE uses this metric to route the tunnel instead of allowing the router itself to choose a path.</p>

Table 123: Fields on the Add Anycast Group Tunnel Page (*Continued*)

Field	Description
Routing Method	<p>From the list, select a routing method for the tunnel to specify whether the PCE should compute and provision the path for the tunnel or not:</p> <p>The available options are:</p> <ul style="list-style-type: none"> <p>routeByDevice—This is the default routing method when a PCC-controlled tunnel is created or learned by the PCE. For this method, The PCE does not compute and provision a path.</p> <p>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.</p> <p>Other routing methods (default, delay, adminWeight, constant, distance, ISIS, OSPF)—When a PCC-controlled tunnel has a routing method that is not routeByDevice, the PCE computes and provisions the path as a strict explicit route when provisioning the tunnel. The tunnel's existing explicit route might be modified to a PCE-computed strict explicit route.</p> <p>For example, a loose explicit route specified by you or learned from the router would be modified to a strict explicit route.</p>
Binding SID	<p>NOTE: This field is available only for segment routing tunnels with NETCONF as the provisioning type.</p> <p>Specify the numerical binding SID label value.</p> <p>This value then becomes the label that represents the path defined by the hops you specify on the Path tab (which are the hops that make up the private forwarding adjacency link).</p> <p>Range: 1000000 to 1048575.</p>

Constraints

Table 123: Fields on the Add Anycast Group Tunnel Page (*Continued*)

Field	Description
Admin Group Include All	<p>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.</p> <p>The admin group bits are mapped to meaningful names, such as colors (configured from the Configuration > Network > Admin Group page), so that you can easily differentiate the different traffic routes in the display and also use coloring constraints to influence the path of the tunnel.</p>
Admin Group Include Any	<p>From the list, select one or more admin group bits for the tunnel to traverse links that include at least one of the admin groups specified in this field. The maximum selections allowed is 32.</p>
Admin Group Exclude	<p>From the list, select one or more admin group bits for the tunnel to traverse links that do not include any of the admin groups specified in this field. The maximum selections allowed is 32.</p>
Maximum Delay	<p>Specify the maximum delay (in milliseconds) for the tunnel, which is used as a constraint for tunnel rerouting.</p>
Maximum Hop	<p>Specify an integer value for the maximum number of hops that the tunnel can traverse.</p>
Maximum Cost	<p>Specify an integer value for the maximum cost to be associated with the tunnel.</p>
<i>Advanced</i>	
Count	<p>Specify the number of parallel tunnels to be created between two endpoints.</p> <p>These tunnels share the same design parameters as specified in the Constraints tab.</p> <p>NOTE: Creating parallel tunnels in this manner is different from provisioning multiple tunnels (Provisioning > Multiple Tunnels) where you configure Design parameters separately for each tunnel.</p>

Table 123: Fields on the Add Anycast Group Tunnel Page (Continued)

Field	Description
Delimiter	<p>NOTE: This field is available only when the count value is greater than 1.</p> <p>Specify a delimiter value, which can consist of alphanumeric characters and special characters except space, comma (,), and semicolon (;).</p> <p>This value is used in the automatic naming of parallel tunnels that share the same design parameters. The PCE names the tunnels using the name you enter in the Properties tab and appends the delimiter value plus a unique numerical value beginning with 1</p> <p>Example: myTunnel_1, myTunnel_2, and so on.</p>
Description	Specify a comment or description for the tunnel for your reference.
IP Z	<p>From the list, select the IP address for Node Z (that is, the egress node).</p> <p>The options in the list are populated based on the Node Z that you selected in the <i>Properties</i> tab in this page.</p>
Symmetric Pair Group	<p>Specify a unique name for the symmetric pair group. You can use any number of alphanumeric and special characters.</p> <p>Tunnels with the same group name (as specified in this field) are considered part of a symmetric pair group.</p> <p>You create a symmetric pair group so that the tunnel from the ingress node to the egress node follows the same path as the tunnel from the egress node to the ingress node. When there are two tunnels with the same end nodes but in opposite directions, the path routing uses the same set of links. For example, suppose Tunnel1 source to destination is NodeA to NodeZ, and Tunnel2 source to destination is NodeZ to NodeA. Selecting Tunnel1-Tunnel2 as a symmetric pair group places both tunnels along the same set of links. Tunnels in the same group are paired based on the source and destination node.</p>
Create Symmetric Pair	<p>NOTE: This option is displayed only when you specify a symmetric pair group.</p> <p>Click the toggle button to enable the creation of a symmetric pair.</p> <p>This option allows you to create the symmetric pair in the same operation as creating the tunnel.</p>

Table 123: Fields on the Add Anycast Group Tunnel Page (Continued)

Field	Description
Diversity Group	Specify the name of a group of tunnels to which this tunnel belongs, and for which diverse paths are desired.
Diversity Level	<p>From the list, select the level of diversity for the tunnel:</p> <ul style="list-style-type: none"> • Default—No diversity level will be applied. • 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. • 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. • Link—Two paths don't intersect at any given link. Link diversity is the weakest.
Slice Include All	Specify one or more topology slice IDs for the tunnel to be routed over links and nodes, that are tagged with all the slice IDs specified in this field.
Slice Include Any	Specify one or more topology slice IDs for the tunnel to be routed over links and nodes, that are tagged with atleast one of the slice IDs specified in this field.
Slice Exclude	Specify one or more topology slice IDs for the tunnel to be routed over links and nodes, that are not tagged with any of the slice IDs specified in this field.
Route on Protected IP Link	Click to enable the toggle button if you want the route to use protected IP links as much a possible.
Custom Attributes	<p>Click the Add icon (+) to specify provisioning properties not directly supported by the GUI.</p> <p>For example, you cannot specify a hop-limit when you provision a tunnel. However, you can add hop-limit as a custom attribute.</p> <p>At the edit > protocols > mpls > label-switched-path hierarchy level in the NETCONF template file, you must add the statements needed to provision with the property you are adding. If the property is present with the defined value, then the provisioning statement is executed.</p>

Table 123: Fields on the Add Anycast Group Tunnel Page (*Continued*)

Field	Description
<i>Path</i>	
Routing Path Type	<p>From the list, select the type of routing path for the tunnel:</p> <ul style="list-style-type: none"> • 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.
Add Hop	<p>This option available only if the routing path type is Preferred or Required.</p> <p>Click the Add (+) icon or click Add Hop. From the list, select an option as the first hop between node A and node Z.</p> <p>In addition, click the toggle button next to this field to specify whether the hop is strict or loose:</p> <ul style="list-style-type: none"> • 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. <p>To add additional hops, click the + icon again. You can add a maximum of 37 hops.</p> <p>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.</p>
<i>Schedule</i>	

Table 123: Fields on the Add Anycast Group Tunnel Page (Continued)

Field	Description
Plan	<ul style="list-style-type: none"> • No Schedule—(Default) tunnel provisioning is not scheduled (that is, tunnels are provisioned immediately upon submission of the provisioning request). • Once—In the Start and End fields that appear, specify the start date and time and end date and time at which you want to provision the tunnels. The tunnels are provisioned once at the specified date and time. • Recurring Daily—Specify the start and end dates and start and end times in the Start Date, End Date, Start Time, and End Time fields that appear. The tunnels are provisioned daily.

Tunnels

IN THIS CHAPTER

- [Understand LSP Delegation and Undelegation | 815](#)
- [Add and Remove LSP Delegation | 816](#)
- [About the Bandwidth Calendar Page | 819](#)
- [About the Path Optimization Page | 821](#)

Understand LSP Delegation and Undelegation

IN THIS SECTION

- [Understand the Behavior of Delegated and Undelegated LSPs | 816](#)

In Paragon Pathfinder, the Path Computation Element (PCE) computes paths in the network and applies computational constraints.

You can delegate the management of Path Computation Client-controlled LSPs (PCC-controlled LSPs) to the PCE either by using the Paragon Automation GUI (from the Configure LSP Delegation page [**Network > Tunnels > Configure LSP Delegation**]) or the Junos OS CLI. After delegation, such LSPs are managed by the PCE and are classified as PCC-delegated LSPs.

When LSPs are managed by the PCE, you can reprovision LSPs (directly from the GUI) for which provisioning has failed or for which the path isn't the expected path. In addition, you can also enable optimization to re-establish an optimal set of paths for the network. Pathfinder does not support reprovisioning and optimization for PCC-controlled LSPs.

You can also delegate the LSPs back to the PCC (also known as undelegating the LSPs from the PCE), after which the LSPs are managed by the PCC and are re-classified as PCC-controlled LSPs. This is useful if Pathfinder is down or if the LSPs are not working as expected.

For the procedure to delegate and undelegate LSPs from the Paragon Automation GUI, see ["Add and Remove LSP Delegation"](#) on page 816.

Understand the Behavior of Delegated and Undelegated LSPs

In both standalone (where the Paragon Automation components run on a single primary node) and high availability (HA) cluster configurations (where the Paragon Automation 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 Pathfinder 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 Pathfinder 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.

RELATED DOCUMENTATION

| [Understand How Pathfinder Handles LSPs](#) | 715

Add and Remove LSP Delegation

IN THIS SECTION

● [Add LSP Delegation](#) | 817

● [Remove LSP Delegation](#) | 818

You can use the Configure LSP Delegation page (**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.

Add LSP Delegation

To delegate one or more LSPs to the PCE:

NOTE: For LSPs that were manually created by using the CLI of IOS-XR devices, you must run device collection before doing any LSP delegation. See ["Add a Device Collection Task" on page 978](#).

1. 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.

2. 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

NOTE: Undelegating LSPs from the PCE is not the same as the temporary removal you achieve when you right-click a tunnel in the Tunnel tab of the network information table and select **Return Delegation to PCC**. In that case, control is temporarily returned to the PCC for a period of time based on the router's timer statement.

To undelegate one or more LSPs from the PCE:

1. 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.

2. 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.

RELATED DOCUMENTATION

| [About the Tunnel Tab](#) | 710

About the Bandwidth Calendar Page

IN THIS SECTION

- [Tasks you Can Perform | 819](#)
- [Field Descriptions | 820](#)

To access this page, select **Network > Traffic Engineering > Bandwidth Calendar**.

Use the Bandwidth Calendar page to view a calendar that displays the LSPs scheduled for provisioning. You can also view the LSP details, provisioning schedule, and bandwidth allocated for the LSPs.

Tasks you Can Perform

- View the details of scheduled LSPs—On the calendar that is displayed on the left side of the page, the following views are supported:
 - Month view (default)—The Month view displays the number of LSPs scheduled to be provisioned on a particular day of the month on the top of that day in the calendar.
 - Year view—The Year view displays the number of LSPs scheduled to be provisioned in a particular month of the year on the top of that month in the calendar.

Click the day (in Month view) or month (in Year view) to view the details of these LSPs in the Scheduled LSPs table on the right side of the page.

You can also click **Show All** on the top-right corner of the table to view all the scheduled LSPs.

NOTE: The Scheduled LSPs table is empty until one or more LSPs are scheduled for provisioning.

- Show or hide columns displayed in the Scheduled LSPs table—Click the **Vertical Ellipsis** icon at the top right corner of the page and select the columns that you want displayed in the Scheduled LSPs table.

Only the columns that you selected are displayed.

- Update the list of scheduled LSPs—Click the **Refresh** icon (circular arrow) at the bottom of the Scheduled LSPs table to refresh (update) the list of scheduled LSPs.

If there's any discrepancy between the LSPs scheduled for provisioning and the LSPs displayed here, clicking this icon will display the latest data.

Field Descriptions

[Table 124 on page 820](#) describes the fields on the Bandwidth Calendar page. The values for these fields are populated based on what you configured when adding the LSPs.

Table 124: Fields on the Bandwidth Calendar Page

Field	Description
Name	Name of the LSP that is scheduled for provisioning.
From	IP address of the ingress node (Node A).
To	IP address of the egress node (Node Z).
LSP Index	Automatically assigned numerical value to identify the LSP in the network model.
Bandwidth	Planned bandwidth configured for the LSP.
Start	Date (in MM:DD:YYYY format) and time (in HH:MM:SS 12-hour or AM/PM format) at which the LSP provisioning starts.
End	Date (in MM:DD:YYYY format) and time (in HH:MM:SS 12-hour or AM/PM format) at which the LSP provisioning ends.
Repeats	Interval at which the provisioning task repeats. Displays — if the LSP provisioning is scheduled only once and Daily if the LSP provisioning is scheduled to recur daily.
Recurrence Ends	Date (in MM:DD:YYYY format) and time (in HH:MM:SS 12-hour or AM/PM format) at which the recurrence ends (if the LSP provisioning is scheduled to recur daily) and — if the LSP provisioning is scheduled only once.
Path Type	Path type for the LSP (Primary , Secondary , or Standby).

Table 124: Fields on the Bandwidth Calendar Page (Continued)

Field	Description
Control Type	Delegated (for PCC-delegated LSPs), PCC (for PCC-controlled LSPs), PCE-initiated (for PCE-initiated LSPs), and Transport (for transport LSPs).
Provisioning Type	Provisioning type for the LSP (RSVP or SR (which indicates segment routing)).
Operational Status	Operational status (Up , Down , or Active) of the LSP.
Last Status	Last status message received from the PCS when processing the LSP.

RELATED DOCUMENTATION

[Add a Single Tunnel | 729](#)

[Add Diverse Tunnels | 743](#)

[Add Multiple Tunnels | 754](#)

About the Path Optimization Page**IN THIS SECTION**

- [Tasks You Can Perform | 822](#)

To access this page, select **Network > > Traffic Engineering > Path Optimization**.

You can use the Path Optimization page to run path analysis and generate a report to help you determine whether optimization should be done. Based on the analysis, you can choose to optimize paths on demand from this page or set a timer to trigger the optimization automatically at regular intervals from the Pathfinder page (**Configuration > Network > Pathfinder**).

NOTE: Path analysis and optimization are applicable only to Path Computation Element-initiated LSPs (PCE-initiated LSPs) and Path Computation Client-delegated LSPs (PCC-delegated LSPs) because the PCE doesn't attempt to optimize PCC-controlled LSPs.

Tasks You Can Perform

- Analyze the network—You can analyze the network for optimization opportunities. The Path Analysis report is generated automatically after path analysis and you can use this report to determine whether optimization should be done.

To trigger path analysis, click **Analyze** in the top-right corner of the Path Analysis table.

All the PCE-initiated and PCC-delegated LSPs in the network are analyzed automatically for optimization opportunities, and the Path Analysis report is generated.

This report lists the LSPs that are currently not in an optimized path, suggests what the optimized paths should be, and provides data about what could be gained (in terms of the number of hops, admin weight, and so on) if the LSPs were to be optimized.

- Optimize paths in the network—Optimization enables the Path Computation Element (PCE) to re-establish an optimal set of paths for a network by finding the optimal placement of LSPs using the current set of nodes and links in the network.

To trigger path optimization, click **Optimize** in the top-right corner of the Path Analysis table.

An alert message appears asking you to confirm the operation.

If you click **Yes**, the PCE sends a provisioning order to the network to optimize all PCE-initiated and PCC-delegated LSPs. In case of provisioning delays or recent changes in the network, low priority LSPs or low bandwidth LSPs may be temporarily unrouted. The PCE will attempt to reprovision these unrouted LSPs until it finds a path.

After optimization is complete, the following reports are generated:

- Path Optimization—Lists the optimized LSPs and their details (such as the path name and optimized path).
- LSP Path Changes—Lists changes (in terms of delay, number of hops, path cost, and so on) to the LSPs as a result of optimization.
- RSVP Link Changes—Lists changes in link RSVP bandwidth reservation (in terms of bandwidth, utilization, and so on) for the LSPs as a result of optimization.

NOTE: The optimization is based on the current network conditions, not on the conditions in effect the last time the analysis was done.

- View the most recent reports—You can view the most-recently generated Path Analysis, Path Optimization, LSP Path Changes, and RSVP Link Changes reports.

To view these reports, click **Reports** in the top-right corner of the Path Analysis table and select the type of report that you want to view.

The selected report is displayed on the Path Optimization page.

- Show or hide columns that contain information about the reports—Click the vertical ellipsis icon on the top-right corner of a report and select the check boxes corresponding to the columns that you want to view in the report. Only the selected columns are displayed in the report.

RELATED DOCUMENTATION

| [About the Tunnel Tab](#) | 710

Change Control Management

IN THIS CHAPTER

- [Change Control Management Overview | 824](#)
- [Change Request Workflow | 826](#)
- [About the Change Control Management Page | 829](#)

Change Control Management Overview

IN THIS SECTION

- [Permissions in the Change Control Management System | 825](#)

Change control management enables you to authorize and track change requests related to tunnels (additions, deletions, and modifications).

An administrator or a user with the required privileges must approve and deploy the change requests for the changes to be implemented. After the requests are implemented, you can view the tunnels (in case of additions) and the related updates (in case of modifications or deletions) in the network information table and in the topology map.

The lifecycle of a change request involves the following tasks:

1. Submit a request
2. Approve or reject the request
3. Deploy the request, if approved
4. Archive the request, if rejected, approved, or activated

For more details, see ["Change Request Workflow" on page 826](#).

You can monitor the status of change requests from the Change Control Management page (**Network > Change Control Management**).

Permissions in the Change Control Management System

What you can do within the change control management system depends on your assigned user role. Each user role has specific permissions associated with it, enabling users with that role to perform various tasks. You can assign the permissions for each user role from the Roles page (**Administration > User Management > Roles**).

[Table 125 on page 825](#) explains the permissions in detail.

Table 125: Permissions in the Change Control Management System

Permission	Description
Create	<p>You can create, modify, and delete change requests, and then submit the change requests to the change control management system.</p> <p>First, you must access the relevant page on the GUI to perform the required task. After you complete the task, a change request is automatically created for the task.</p>
Approve	<p>You can approve or reject change requests that are created by anyone, including those that you created.</p> <p>BEST PRACTICE: Monitor the Change Control Management page regularly and advance change requests promptly to keep the requests moving through the change control management system.</p>

Table 125: Permissions in the Change Control Management System (*Continued*)

Permission	Description
Auto Approve	<p>You can create change requests that are automatically approved and deployed. The Auto-approve permission option also applies to change requests that are made by using REST APIs, thus enabling automated northbound integration with third-party systems or scripts.</p> <p>NOTE:</p> <ul style="list-style-type: none"> • The Create option and the Auto-approve permission options are mutually exclusive because the Auto-approve permission option includes the Create option. • When you deploy a request without selecting the Auto-approve permission option, you can schedule the provisioning for a later date and time. However, with the Auto-approve permission option, the approval and deployment are immediate, bypassing the scheduling step. • The Auto-approve permission option doesn't enable you to approve change requests that are submitted by other users.
Deploy	<p>You can deploy approved change requests that are created by anyone, including those that you created.</p> <p>BEST PRACTICE: Monitor the Change Control Management page regularly and advance change requests promptly to keep the requests moving through the change control management system.</p>
If you have none of these permissions	You can only view the status of the change requests.

RELATED DOCUMENTATION

[About the Change Control Management Page | 829](#)

[About the Tunnel Tab | 710](#)

Change Request Workflow

Change control management provides authorization and tracking for change requests related to tunnels.

The following procedure explains the workflow of a change request. Based on your user role, you can perform one or more of the following steps:

1. Create and submit a change request.

NOTE: You can create a change request only if you have either the Create or Auto-approve permission assigned to your user role.

To create and submit a change request:

- a. Navigate to the Tunnel tab in the network information table (**Network > Topology**).
- b. Perform the task (add, modify, or delete tunnels) appropriate for the desired request. After you complete the task, a change request is automatically created for the task and is submitted to the change control management system. The request appears on the Change Control Management page (**Network > Change Control Management**) with Status as *Submitted*.

If you have the Auto-approve permission assigned to your user role, the requests are approved and deployed automatically upon submission. If you don't have this permission, you or a user having the Approve permission must approve the request. See [Step 2](#).

2. Approve or reject a change request.

NOTE: You can approve or reject the change request only if you have the Approve permission assigned to your user role.

To approve or reject the change request:

- a. On the Change Control Management page, select the request that you want to approve and click **Approve**. Alternatively, click **Reject** to reject the request. The Approve Change Request *request-id* page or Reject Change Request *request-id* page appears, based on what you selected in Step 1.
- b. (Optional) Add a comment and click **OK**. A confirmation message appears on top of the page, and the status of the request changes accordingly.

3. Deploy the approved change request on the network for the request to take effect.

NOTE: You can deploy the change request only if you have the Deploy permission assigned to your user role.

To deploy the change request:

- a. Select the request that you want to deploy, and click **Deploy**.
The Deploy Change Request *request-id* page appears.

NOTE: This step is not applicable when change requests are auto-approved. Auto-approved change requests are approved and deployed automatically upon submission.

- b. You can choose to start deployment immediately (**Activate Now**) or schedule the deployment for a later date and time (**Activate Later**).
- c. (Optional) Specify a comment for the request.
- d. Click **Deploy**.

After a change request is scheduled, the status of the request changes to *Scheduled*.

When you deploy this request at the scheduled time, the status changes to *Activated*. The Path Computation Element (PCE) provisions the tunnel and the tunnel appears in the network information table (Tunnel tab) in the Topology page.

4. (Recommended) Archive (close) the approved, activated, or rejected change requests when they are no longer needed.

NOTE: You can archive a request only if *you* have submitted it. You cannot archive requests that are submitted by other users.

To archive a request:

- a. Select a request with the Status *Rejected*, *Approved*, or *Activated*, and click **More > Archive**.
The Archive Change Request *request-id* page appears.
- b. (Optional) Specify a comment in the Approver Comments field.
- c. Click **Archive**.

A confirmation message appears on top of the page, indicating that the change request is closed. You can view this request (Status: *Closed*) in the Archived tab on this page.

RELATED DOCUMENTATION

[Change Control Management Overview](#) | 824

About the Change Control Management Page

IN THIS SECTION

- [Tasks You Can Perform | 829](#)
- [Field Descriptions | 830](#)

To access this page, select **Network > Change Control Management**.

Use the Change Control Management page to view and manage change requests and archived requests in the change control management system.

Tasks You Can Perform

- View submitted change requests in the **Change Requests** tab and archived requests in the **Archived** tab. To view details of a request, click the **Details** icon (that appears when you hover over the request). Alternatively, select the request and click **More > View Details**. You can also click the Toggle Details icon that appears next to the Custom icon (vertical ellipsis).

The Change Request *request-ID* pane appears on the right side of the page, displaying the details of the request.

- Approve or reject change requests—If you are assigned the Approve permission, you can approve or reject one or more change requests created by anyone, including those created by yourself. See ["Change Request Workflow" on page 826](#).
- Deploy the approved change requests—If you are assigned the Deploy permission, you can deploy one or more approved change requests created by anyone, including those created by yourself. See ["Change Request Workflow" on page 826](#).
- Archive change requests—You can archive one or more change requests (created by you) with the status *Submitted*, *Rejected*, or *Activated*. See ["Change Request Workflow" on page 826](#).

NOTE: You can archive a request only if *you* have submitted it. You cannot archive requests that have been submitted by other users.

- Export change requests and archived requests—From the respective tabs, click **More > Export All** to export change requests or archived requests.

The requests are downloaded to your system as a comma-separated values (CSV) file.

- **Modify submitter comments**—Select a change request that you've submitted and click **More > Modify Submitter Comment**.

In the Modify Submitter Comment page that appears, specify your comment, and click **OK**.

The updated comment appears in the Submitter Comment column.

- **Reschedule deployment of change requests**—Select one or more change requests for which you want to reschedule deployment, and click **More > Reschedule**.

The Reschedule Change Request *request-id* page appears, where you can choose to specify a comment. Then, click **Reschedule**.

A confirmation message appears; the status of the rescheduled requests reverts to *Approved*.

NOTE: You can reschedule only requests that are in the **Scheduled** state.

- **Show or hide columns displayed on the page**—Click the **Custom** icon (vertical ellipsis) in the top-right corner of the page and select the columns that you'd like to see on the page. Only the selected columns are displayed.

Field Descriptions

Table 126 on page 830 describes the fields on the Change Control Management page.

Table 126: Fields on the Change Control Management page

Field	Description
Status	Current status of the change request (Submitted, Approved, Rejected, Scheduled, or Activated).
ID	ID that is automatically generated for the change request.
Request Action	Action that is expected from the change request (Add, Modify, or Delete).
Submitter	Name of the user who has submitted the change request.

Table 126: Fields on the Change Control Management page (*Continued*)

Field	Description
Submitted Time	Date (in MM:DD: YYYY format) and time (in HH:MM:SS 12-hour or AM/PM format) at which the change request was submitted for approval.
Submitter Comment	Automatically-generated comment that reflects the action (such as Add LSP) for change requests. If you've modified the comment from the Modify Submitter Comment page (More > Modify Submitter Comment), the modified comment is displayed here.
Approver	Name of the user who has approved the change request.
Approver Time	Date (in MM:DD: YYYY format) and time (in HH:MM:SS 12-hour or AM/PM format) at which the change request was approved.
Approved Comment	Comment provided by the user who has approved the change request. For auto-approved requests, this column displays Auto Approved .
Type	Type of change request (LSP).
Provisioning Status	Provisioning status (such as Activation Scheduled, Successful) of the LSP and errors (such as Unable to add planned LSP) based on the status of the change request. Example: If the LSP is created successfully, the status is displayed as Successful .

RELATED DOCUMENTATION

[Change Control Management Overview](#) | 824

9

PART

Monitoring

[Monitor Network Health | 833](#)

[Manage Alarms and Alerts | 845](#)

[Monitor Jobs | 854](#)

[Analytics | 858](#)

[Monitor Workflows | 899](#)

Monitor Network Health

IN THIS CHAPTER

- [About the Network Health Page | 833](#)
- [Activate Time Inspector View | 844](#)

About the Network Health Page

IN THIS SECTION

- [Timeline View | 834](#)
- [Tile View | 836](#)
- [Table View | 840](#)
- [Time Inspector View | 842](#)

Paragon Automation offers a way to visualize device-level and network-level health problems through the Network Health page. You can visualize network health data in a time line — of devices, device groups, and network groups — or at a more granular levels of keys so that you can discover the root cause of issues detected by the platform.

Use the Health page (**Monitoring > Network Health**) to monitor and track the health of a single device, a device group, or a network group. You can also troubleshoot problems. Select an entity type (DEVICE, DEVICE GROUP, or NETWORK) located in the top left corner of the page. Once you select an entity type, you can then select a device name, device group name, or a network group name from the drop-down menu. From the date and time displayed below the *Entity Type*, you can select any date and time (in AM/PM format) for which you want to view the network health data.

You can click **Save as Default** after selecting a device or device group to save that setting as the default view on the Network Health page. Click **Clear Default** to clear the saved settings.

The page is divided into the following three main views that, when used together, can help you investigate the root cause of problems detected on your devices:

- ["Timeline View" on page 834](#)
- ["Tile View" on page 836](#)
- ["Table View" on page 840](#)

You can navigate to the ["Time Inspector View" on page 842](#) from the Table View, if you chose device group as the entity type and selected multiple devices.

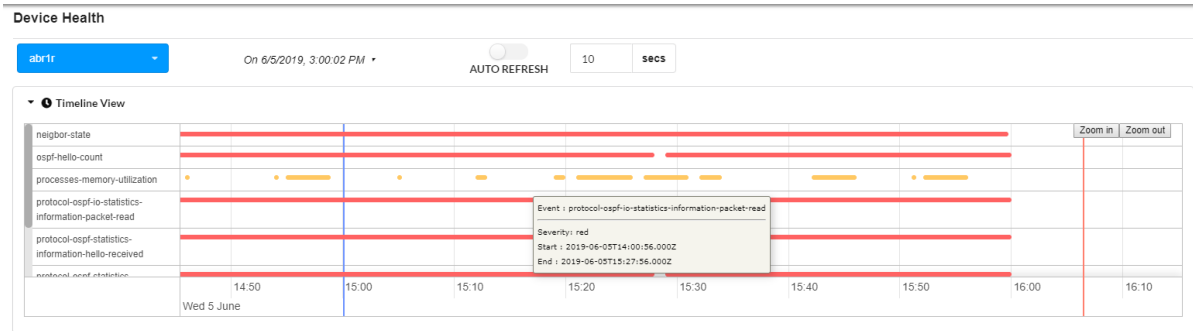
Timeline View

In timeline view, you can monitor real-time and past occurrences of KPI events flagged with a minor or major severity level health status. The general characteristics and behaviors of the time line include (see [Figure 64 on page 835](#)):

- Clicking on the right caret next to the Timeline View heading expands or collapses the time line.
- Each dot or line in the time line represents the health status of a unique KPI event (also known as a Paragon Insights rule trigger) for a pre-defined KPI key with which Paragon Insights has detected a minor or major severity level issue. The name of each event is displayed (per device) directly to the left of its associated health status dot or line.
- The health status dot or line for each unique KPI event in the time line can consist of several different KPI keys. Use tile view and table view to see the health status information for the KPI keys.
- Only minor or major severity level KPI events are displayed in the time line. Yellow represents a minor event and red represents a major one.
- A KPI event that occurs once (at only one point in time) and does not recur continuously over time is represented as a dot.
- A KPI event that occurs continuously over time is represented as a horizontal line.
- Time line data is by default displayed for a two-hour customizable time range. However, you can enlarge the time line graph (shorten the time range covered) by clicking on the **Zoom in** button and shrink the time line graph (extend the time interval covered) using the **Zoom out** button at the top right corner of the graph.
- The red vertical line on the time line represents the current time.
- The blue vertical line on the time line represents the user-defined point of time for which to display data. If you double click on any other time point in the time line graph, the blue line shifts to the location where you clicked and the graph displays two hours before and after the current position of the blue vertical line.

The Tile View and the Table View are updated to populate data on trigger evaluations and key performance indicators for the current time represented by the blue line.

Figure 64: Timeline view



The following table describes the main features of the time line:

Feature	Description
Display information about a dot or horizontal line in the time line.	Hover over the dot or horizontal line to display the associated KPI event name, device name, health status severity level, and event start and end times. Additional health status information about the KPI event can be found in tile view. For information about tile view, see the "Tile View" on page 836 section.
For the displayed data, change the range of time (x-axis) that is visible on the page.	Options: <ul style="list-style-type: none"> • Click and drag the x-axis of the time line to the left or to the right. • Click the Zoom In or Zoom Out buttons in the top right corner of the time line.

(Continued)

Feature	Description
Choose a different two-hour time range of data to display.	<p>Use the blue vertical line to customize the time range of data to display. Options for enabling the blue vertical line:</p> <ul style="list-style-type: none"> • Click inside the time line grid at the particular point in time you want to display data. • In the date/time drop-down menu (located above the time line), select the particular point in time you want to display data. <p>Data is generally displayed for 1 hour before and 1 hour after the blue line. Hover over the blue line to display the exact point in time that it represents. Drag the blue line left or right to adjust the time.</p> <p>NOTE: Auto-refresh is disabled whenever you enable the blue line. Re-enabling auto-refresh disables the blue line and resets the time line to display the most recent two-hour time range of data.</p>
Freeze the time line (disable auto-refresh).	Toggle the auto-refresh switch to the left.
Unfreeze the time line (enable auto-refresh).	Toggle the auto-refresh switch to the right.

Tile View

The tile view uses colored tiles to allow you to monitor and troubleshoot the health of a device. The tiles are organized first by topic and by unique KPI key (see [Figure 66 on page 838](#)) for the device or for the devices you selected in a device group or network group. By default, the tile view data corresponds to the most recent data collected. To customize the point in time for which data is displayed in tile view, select a particular point in time from the date/time drop-down menu (located above the time line) or enable the blue vertical line in timeline view. For information about how to enable the blue vertical line, see the ["Timeline View" on page 834](#) section.

Figure 65: Tile View of a Device Group

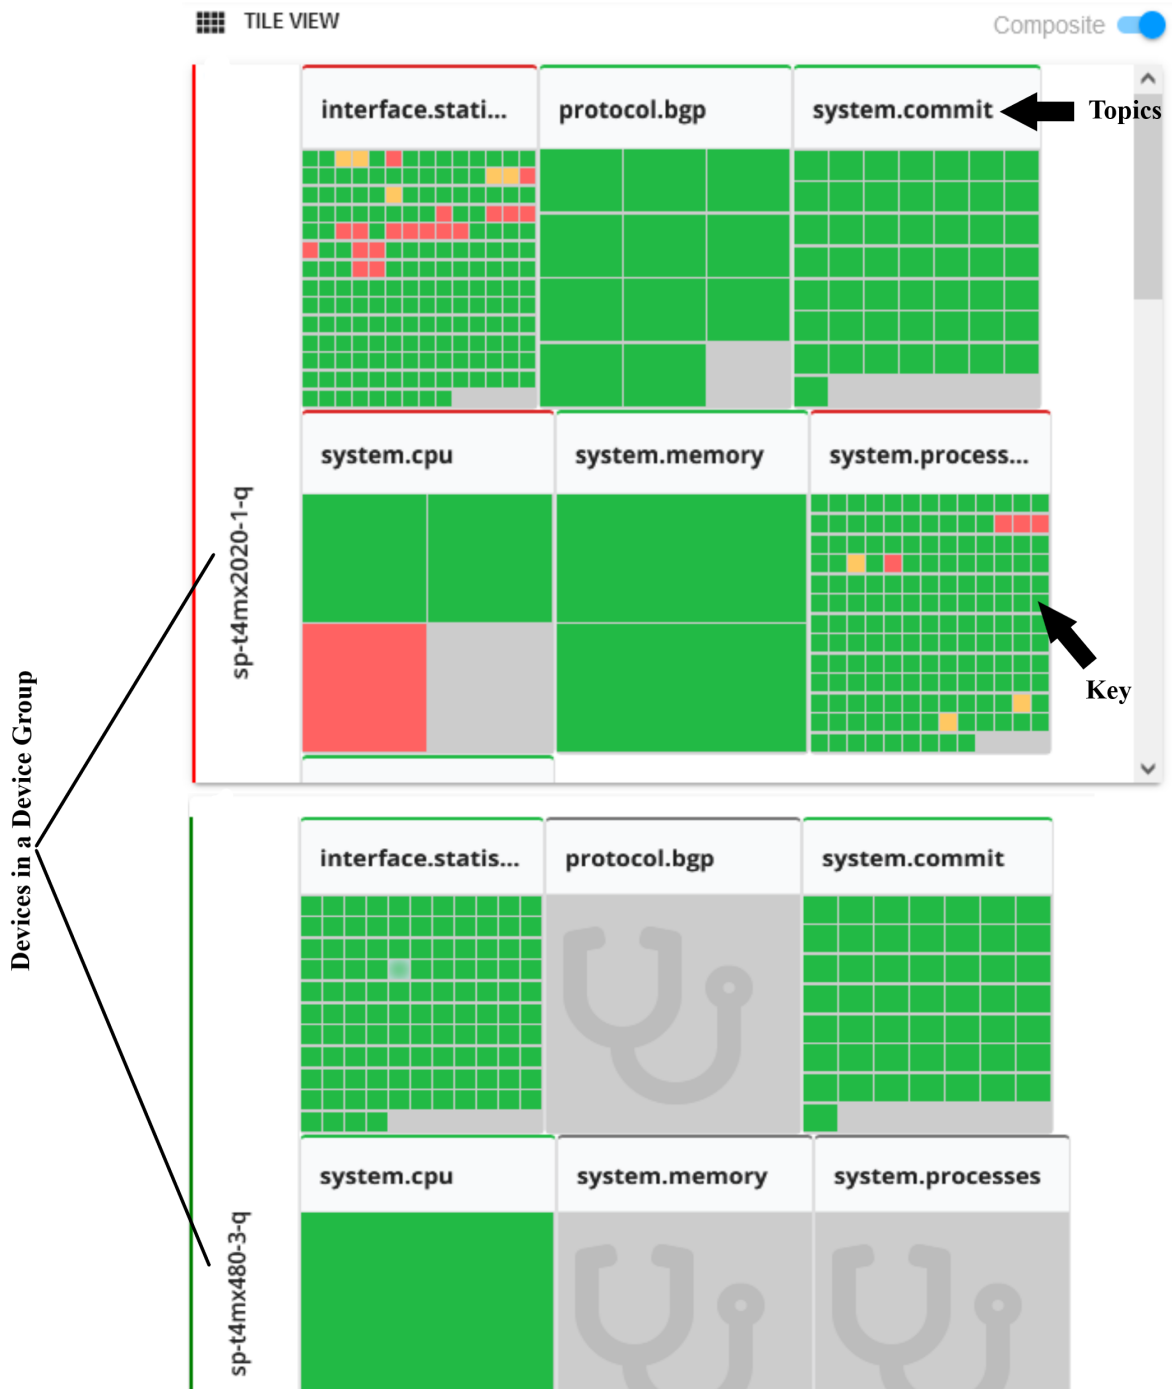



Figure 66: Tile View of a Device



The following table describes the meaning of the severity level colors displayed by the status tiles:

Color	Definition
Green	The overall health of the KPI key is normal. No problems have been detected.
Yellow	There might be a problem with the health of a KPI key. A minor problem has been detected. Further investigation is required.
Red	The health of a KPI key is severe. A major problem has been detected.
Gray	No data is available.

The following table describes the main features of the tile view:

Feature	Description
Display information about a status tile.	<p>Options:</p> <ul style="list-style-type: none"> • Hover over a status tile to display the name of the key, KPIs associated with the key, and the status messages associated with the KPIs. • Click on a status tile. Information about the status tile is displayed in table view. For information about table view, see the "Table View" on page 840 section. <p>Note: If the number of KPI keys exceeds 220, the keys are automatically aggregated and grouped.</p>
Display information in table view about the status tiles associated with a single topic.	Click on a topic name in tile view. For information about table view, see the "Table View" on page 840 section.
Composite Toggle	When active, users can click on specific keys within the tile groups. This allows you to pass multiple KPIs to the Time Inspector View.

Table View

The table view allows you to monitor and troubleshoot the health of a single device provided in a customizable table. You can search, sort, and filter the table data to find specific KPI information, which can be especially useful for large network deployments. To select which attributes are displayed in the table, check the appropriate check box in the field selection bar above the table (see [Figure 67 on page 840](#)). The checkbox on the left side of each row is used to help activate the Time Inspector view. Multiple rows can be selected at one time.

Figure 67: Table View

The screenshot shows the 'TABLE VIEW' interface. At the top right, there is a 'TIME INSPECTOR' button. Below it is a 'Field Selection Bar' with checkboxes for Time, Device, Group, Topic (checked), Rule, Keys (checked), KPI (checked), Status (checked), and Message (checked). Below the selection bar is a table with columns: Topic, Keys, KPI, Status, and Message. The 'Topic' column has a search field with 'system.stor...' entered. The 'Status' column has a dropdown menu set to 'ALL'. The table contains five rows of data, each with a checkbox on the left. The first row has a green status bar, and the others have red status bars. The messages describe storage usage for various file systems.

The following table describes the attributes supported in table view:

Attributes	Description
Time	Time and date the event occurred.
Device	Device name.
Group	Device group name.
Topic	Rule topic name.
Keys	Unique KPI key name.

(Continued)

Attributes	Description
KPI	Key Performance Indicator (KPI) name associated with an event.
Status	Health status color. Each color represents a different severity level.
Message	Health status message.

The following table describes the meaning of the severity level colors displayed by the **Status** column:

Color	Definition
Green	The overall health of the KPI key is normal. No problems have been detected.
Yellow	There might be a problem with the health of a KPI key. A minor problem has been detected. Further investigation is required.
Red	The health of a KPI key is severe. A major problem has been detected.

The following table describes the main features of the table view:

Feature	Description
Sort the data by ascending or descending order based on a specific data type.	Click on the name of the data type at the top of the column by which you want to sort.
Filter the data in the table based on a keyword.	Enter the keyword in the text box under the name of a data type at the top of the table (see Figure 67 on page 840).

(Continued)

Feature	Description
Navigate to a different page of the table.	Options: <ul style="list-style-type: none"> • At the bottom of the table, click the Previous or Next buttons. • At the bottom of the table, select the page number using the up/down arrows (or by manually entering the number) and then press Enter.
If the data in a cell is truncated, view all of the data in a cell.	Options: <ul style="list-style-type: none"> • Hover over the cell. • Resize the column width of the cell by dragging the right side of the title cell of the column to the right.
Row selection checkbox	Make this row's data available for Time Inspector view.

Time Inspector View

Time Inspector is a composite view that provides a timeline view of trigger conditions on KPI data that you selected in **Table View**. You can also drag and drop trigger conditions to view the conditions in one graph or as separate graphs. Time Inspector was initially available only when the entity type **DEVICE GROUP** is selected. However, **Time Inspector View** is available when you select **Device** or **Network** entity type as well. After you select an entity type, you can access time inspector view by clicking the **TIME INSPECTOR** button located below the **Timeline View** pull-down.


This view allows you to drill down into field-level data for specific triggers over a time line.

When the **Health** page is first accessed, the **TIME INSPECTOR** button is disabled. For more information on activating the **TIME INSPECTOR** button, see ["Activate Time Inspector View" on page 844](#) .

The **Composite** toggle switch at the upper right of the **TILE VIEW** is used along with Time Inspector View. The **Composite** toggle switch allows you to select data from more than one topic to be shown in the **Table View** and, thus, the Time Inspector View. This can be useful when topics must be combined to find root cause for an issue. For example, system memory usage could combine with output queue usage to create a performance issue in an overloaded system. In certain cases, such as for keys that represent interface statistics on bandwidth utilization, a single tile could represent more than one key.

Figure 68 on page 843 below shows a time inspector window created from the system.storage usage topic for a specific device.

Figure 68: Time Inspector Window



As you can see, the Time Inspector window has a mini time line at the top, a topic-based line chart below, and a chart selector section at the bottom. This particular chart was created as a composite (indicated by the merging blue arrow) of a file-system-utilization in the check-storage rule of the system.storage topic.

Note that there are 3 fields in the check-storage rule: used-percentage, low-threshold, and high-threshold. Since the chart was created as a composite (fields charted together) there are three lines on the displayed chart. If the “chart fields separately” button (diverging arrows) were clicked instead, you would see 3 single-line charts showing the same data.

The more rules you select with the **TABLE VIEW** check boxes, the more charts you can create in the **Time Inspector** view.

RELATED DOCUMENTATION

| [Activate Time Inspector View](#) | 844

Activate Time Inspector View

Time Inspector is a composite view that provides a timeline view of trigger conditions on KPI data that you selected in **Table View**. You can also drag and drop trigger conditions to view the conditions in one graph or as separate graphs. For more information, see "[Time Inspector View](#)" on page 842.

To activate the **TIME INSPECTOR** button and make the view available, you must:

1. Select the **Entity Type** from the **Network Health** page.
Time Inspector View was initially available only when the entity type **DEVICE GROUP** is selected. However, **Time Inspector View** is now available when you select **Device** or **Network** entity type.
2. Select one or more devices, device groups, or networks from the drop-down list next to the entity type that you have selected.
3. Have valid data in at least one device, device group, or network component topic in **TILE VIEW**.

NOTE: Topics showing "no data" will not work for enabling the Time Inspector view.

4. Have data appearing in the **TABLE VIEW** section. You can achieve this by clicking the device, device group, or network component topic header in **TILE VIEW**.
5. Select the checkbox to the left of at least one of the rows in **TABLE VIEW**.

The **TIME INSPECTOR** button is enabled.

When clicked, the **TIME INSPECTOR** button opens the **Time Inspector View** pop-up window. For more information, see "[Time Inspector View](#)" on page 842.

RELATED DOCUMENTATION

| [About the Network Health Page](#) | 833

Manage Alarms and Alerts

IN THIS CHAPTER

- [Alerts and Alarms Overview | 845](#)
- [About the Alarms Page | 847](#)
- [About the Alerts Page | 850](#)

Alerts and Alarms Overview

You can view alarms raised on devices on the Alarms (**Monitoring > Alarms and Alerts > Alarms**) page. An alarm indicates conditions on a device that might prevent the device from operating normally. Alarm conditions for a device are predefined and are raised based on the fault monitoring and performance monitoring (FMPM) being performed on the device. For example, conditions such as hardware issues, drop in throughput and latency of data, temperature variations, and capacity optimization issues automatically trigger an alarm.

You can view alarms raised on both Juniper Networks and Cisco IOS-XR devices. On Juniper Networks devices, you can view alarms for the following conditions:

- Interface down
- Chassis alarms
- High output utilization
- High input utilization
- High CPU utilization
- High memory utilization
- High storage utilization
- Fan failure
- High temperature

For Cisco IOS XR devices running IOS XR Release 7.1.5 and later, you can view alarms defined in the OpenConfig alarm model.

NOTE: For Cisco IOS XR devices, you must set the default NETCONF port to 22, otherwise, you cannot view alarms.

You can view alarm statistics on the Paragon Automation dashboard. The dashboard displays the severity of the different alarm types and the alarms history for the previous one hour, one day, one week, one month, and one year.

You can view alerts that Paragon Automation generates on the Alerts (**Monitoring > Alarms and Alerts > Alerts**) page. Paragon Automation generates an alert for anomalies in a device group or a network group. You must deploy playbooks on device groups and network groups to monitor specific key performance indicators (KPIs) and detect anomalies. Anomalies include system errors, protocol errors, interface errors, chassis errors, and other custom configurable KPIs. Paragon Automation also automatically combines several alerts under a main alert, called smart alarm, identifying the root cause of these combined errors and anomalies. These alerts are known as ["smart alerts" on page 850](#). You can receive an alert notification for major events when you configure a notification profile and enable the notification profile in device group and network group settings. See ["Configure a Notification Profile" on page 621](#) for more information. You can view minor, major, and normal alerts on the Alerts page. You can also track the KPI associated with an alert from the timeline view, tile view, and table view on the Network Health page.

[Table 127 on page 846](#) describes the lifecycles of alarms and alerts.

Table 127: Alarm and Alert Lifecycle

Alarms	Alerts
Raise—An alarm is raised when conditions defined in the telemetry manager are met. Paragon Automation stores the alarms and you can view them on the Alarms page.	Raise—An alert is raised when an anomaly (for example, KPI exceeds a preset threshold) or a status change (for example, a link goes down) is detected on a device group or network group.
Assign—A user is assigned to check the issue raised by the alarm.	Acknowledge—A user views and acknowledges an alert. The user takes steps to remediate the conditions in the device that generated the alert. An acknowledgment indicates that work is in progress to rectify the error or anomaly.

Table 127: Alarm and Alert Lifecycle (*Continued*)

Alarms	Alerts
Acknowledge—A user marks an alarm with the acknowledged status if they have viewed and/or troubleshooting the issue indicated by the alarm.	Shelve— You can set the status of an alert as shelve. Shelve lowers the priority of an alert and snoozes the alert for the time you configure. You can also re-open an alert that is closed or shelved. Or Close—After you resolve the issue raised by an alert, you can set the status of an alert as close. You can also re-open an alert that is closed.
Clear—Paragon Automation automatically clears alarms after the conditions that raised the alarm are normalized.	Delete—To remove an alert from the Alerts page, you can delete the alert.

RELATED DOCUMENTATION

[About the Alarms Page | 847](#)

[About the Alerts Page | 850](#)

About the Alarms Page

Paragon Automation generates alarms on device groups are generated automatically by the Telemetry Manager as part of device discovery.

You can use the Alarm Manager feature to filter, track, and manage alarm notifications received from devices. Alarms alert you to conditions that might prevent the device from operating normally. System alarm conditions are preset. To access the Alarms page, go to **Monitoring > Alarms and Alerts > Alarms**.

In the Alarms page, you can perform the following tasks:

- Assign an alarm — You can assign an alarm to a user. Select an alarm and click the **Assigned** button. The Assign window appears. In the drop-down menu, select a user to whom you want to assign the alarm notifications. Click **OK**.
- Acknowledge the status of an alarm — When a user wants to mark an issue raised by an alarm as seen, you can update the **Acknowledged** status of the issue. Acknowledging an alarm does not clear

the alarm from the page. Select an alarm and click the **Acknowledge** button. The Acknowledge window appears. You can optionally enter an acknowledgement message and click **OK**.

- Filter alarms.

To filter alarms:

1. Click the filter (funnel icon) button.
2. In the drop-down menu, click **Add Filter**.

An Add Criteria window appears.

3. Select the *Field* and *Condition* from the drop-down menus.

Enter a value in the *Value* field. For example, if you want to filter all alarms excluding the chassis alarms, select *Field* as *Source*, *Condition* as *!=*, and enter *Value* as *chassis*.

4. Click the **Add** button.

[Table 128 on page 848](#) describes the attributes in the Alarms page.

Table 128: Alarm Attributes

Attributes	Description
Device	The drop-down menu lists all the devices for which alarms notification is configured. You can select a device from the list.
Severity	<p>You can filter alarms based on the severity level of the alarm. Options include:</p> <ul style="list-style-type: none"> • Major • Minor • All

Table 128: Alarm Attributes (Continued)

Attributes	Description
Time period	<p>In the drop-down menu at the right corner of the Alarms page, you can filter alarms on the basis of time period. The options for time period include:</p> <ul style="list-style-type: none"> • Pre-defined time intervals of 15 minutes, 30 minutes and so on. • Calendar date range – You can use the Select from calendar option to select the start and end date from a calendar. • Relative time – You can select from customized time such as 5 minutes, today, last week, last year and so on. This option is available when you click the More button in the drop-down menu. <p>You can also directly enter a customized time period to filter alarms directly in the time period box.</p> <p>Use the play button near the time range box to apply your filter. If you set the filter for 15 minutes, you can click on the rewind button to apply alarm filter for the previous 15 minutes from the current time in your system.</p> <p>You can use the fast forward button to set the time filter back to current time, if you had used rewind feature.</p> <p>Once you set the time period, click the Apply button to apply your filter.</p>
Time line chart	<p>The time line chart shows the alarms for a selected severity level for the time selected.</p> <p>By default, the severity levels Major and Minor are active. If you click on the severity level once, it would turn the label inactive and hides the chart for the inactive severity level.</p>
Save	<p>The Save button (floppy icon) on the left top corner of the Alarms page, gives you the option to save an alarm query.</p> <p>After you set the time period, you can click on the save drop-down menu, enter your query name and click the Save button. Your query will appear below the default saved queries. To delete your saved query, click on the trash icon that appears beside your saved query.</p>

Table 129 on page 850 describes the fields you see in the Alarms page.

Table 129: Fields on the Alarms page

Field	Description
Severity	Displays an icon representing the severity level – Minor or Major.
Time Raised	Displays the date and time when the alarm was raised in Paragon Insights.
Time Updated	Displays the date and time when the alarm was acknowledged or assigned.
Device	Displays the name of the device affected by the alarm.
Source	Displays the source of alarm such as an interface or system alarm.
Description	Displays details about the alarm in the source. For example, operational status of an interface is down.
Type	Displays the nature of the alarm at the source. For example, a state change (in the interface).
Acknowledged	Displays yes if an alarm is acknowledged.

RELATED DOCUMENTATION

[About the Alerts Page](#) | 850

About the Alerts Page

To access Alerts page, go to **Monitoring>Alarms and Alerts>Alerts**.

You can use the Alerts page to organize, track, and manage KPI event alert notifications received from Paragon Automation devices. Paragon Automation Platform does not track alerts by default. Device groups or network groups are configured to send the alert notifications that are listed in the Alerts page.

Paragon Automation generates smart alarms if you configured resources and dependencies. To configure resources, click **Resource Discovery** at the top right corner of the Alerts page.

Smart alarms combine alarms from different rules into a collapsible tree structure. The main alarm in the tree displays the root cause that triggered the other alerts in the tree. See "[Understand Root Cause Analysis](#)" on page 369 for more information.

Note that Paragon Automation Platform consolidates duplicate alerts into one table entry and provides a count of the number of duplicate alarms it has received.

In the Alerts page, you can enable **Auto Refresh** and enter the time interval (for example, 10s) after which the alerts page must refresh (or update) the listed data. You can also search alerts using the search icon and filter alerts.

To filter alerts:

1. Click the filter (funnel icon) button.
2. In the list, click **Add Filter**.

An Add Criteria window appears.

3. Select the *Field* and *Condition* from the list.

Enter a value in the *Value* field. For example, if you want to filter all alarms excluding the chassis alarms, select *Field* as *Source*, *Condition* as *!=",* and enter *Value* as *chassis*.

4. Click the **Add** button.

The following table describes the alert page attributes.

Table 130: Fields in Alerts Page

Attributes	Description
Severity	Severity level of the alarm. Options include: <ul style="list-style-type: none"> • Major • Minor • Normal
Status	Management status of the alarm entry. Options are Open, Active, Shelved, Closed, and Ack. The statuses available in the Status pull-down menu in the top row of the table only include statuses of alarms visible in the table and those allowed by the status filter above the table.
Last Received	Time the alarm was last received.

Table 130: Fields in Alerts Page (Continued)

Attributes	Description
Dupl.	Duplicate count. Number of times an alarm with the same event, resource, environment, and severity has been triggered.
Topic	Device component topic name.
Resource	Device name.
Event	Name of the rule, trigger or field, and event with which the alarm is associated.
Text	Health status message.

If you click the details button (next to radio button) of an alert, the Alert Details window appears. You can organize alerts by setting the following statuses:

Table 131: Alert Status

Open	You can re-open an alert with <i>closed</i> or <i>shelved</i> status by clicking Open in the Alerts Details window. The alert status changes to <i>open</i> .
Shelve	You can set the status of an alert to <i>shelve</i> if you want to set a lower priority to an alert. You can shelve an alert by one, two, four, or eight hours.
Acknowledge (Ack)	You can set an alert status to <i>ack</i> or <i>acknowledged</i> to mark the alert as seen. The ack status intimates other users to target unacknowledged alerts.
Close	You can filter all alerts with severity level normal and mark their status as <i>closed</i> . An alert with <i>close</i> status is not removed from the Alerts page.
Delete	You can click Delete in the Alert Details window of an alert, to remove the alert.

RELATED DOCUMENTATION

[Enable Alert Notifications for Device Groups and Network Groups | 627](#)

[Configure a Notification Profile | 621](#)

Monitor Jobs

IN THIS CHAPTER

- [About the Jobs Page | 854](#)
- [Viewing Job Details | 856](#)
- [View Job Status | 856](#)

About the Jobs Page

IN THIS SECTION

- [Tasks You Can Perform | 854](#)
- [Field Descriptions | 855](#)
- [Field Descriptions | 855](#)

To access this page, click **Monitoring > Jobs**.

Use this page to view the list of all jobs and the jobs that are scheduled to be executed. You can view general information about the jobs and the overall progress and status of the jobs.

Tasks You Can Perform

You can perform the following tasks from this page:

- View details about a job. See "[Viewing Job Details](#)" on page 856 .
- Sort and filter Jobs:
 - Click a column name to sort the jobs based on the column name.

- Click the **filter** icon and select whether you want to show or hide column filters or apply a quick filter. For example, you can filter jobs based on the Job ID, Start Time, End Time, Owner and so on.

Field Descriptions

[Table 132 on page 855](#) displays the fields on the **Jobs** page.

Table 132: Fields on the Jobs Page

Field	Description
Job Name	The name of the job.
Job ID	The unique identification number assigned to a job.
Status	The status of the job— Success, Failed and In Progress.
Owner	The user who created the job.
Start Time	The time when the job is started or created.
End Time	The time when the job is complete or failed.

Field Descriptions

[Table 133 on page 855](#) provides guidelines on using the fields on the **Scheduled Jobs** page.

Table 133: Fields on the Scheduled Jobs Page

Field	Description
Schedule ID	The uniquely identified number assigned to a scheduled job.
Name	The name of the scheduled job.

Table 133: Fields on the Scheduled Jobs Page (Continued)

Field	Description
Status	Defines the current status of scheduled job. The valid values are Scheduled, Success and Failed.
Owner	View the name of the owner who scheduled the job.
Next Run Time	View the time when the job is scheduled to run next.

Viewing Job Details

To view job details, select a job on the **Jobs** page and click **More>Detail**. The **Details for Job-Name** page appears. This page has the following two tabs:

- The **Details** tab displays the details associated with the job such as Job Name, Status, Start Time, End Time, Owner and Job ID. See **About the Jobs Page** for a description of each of the fields on this page.
- The **Tasks** tab displays the number of tasks associated with the job. A green check mark (success) or a red cross mark (failed) is displayed next to each task indicating the status of the task. Click to list all sub-tasks associated with the task. You can hover on the icon to see the status message and elapsed time information of the task.

Schedule Job will only list those jobs with current status Scheduled or Failed. Once job is scheduled, job status changes to Success , and it won't be listed in **Schedule Job** page.

To quickly view the jobs that are running or that are scheduled, mouse over the **clock** icon, which opens a widget that contains two tabs: In-Progress and Scheduled. The number of jobs that are in progress and scheduled are shown in parentheses in the title of the tab, and each tab lists the in-progress and scheduled jobs. You can view all the jobs by clicking the **See all jobs** hyperlink, which takes you to the **Jobs** page.

View Job Status

View detailed information about an executed job and the status of the corresponding tasks on the Job Status page.

To view the Job Status page, click a job name in the Jobs (**Monitoring > Jobs**) page. The Job Status page appears displaying information about the job. For information about the fields on the page, see Table 1.

You can also view the status and a list of all tasks and sub-tasks in a job at the bottom of the page. If all the sub-tasks in a task are successful, the task has a check mark in a green circle displayed next to it indicating a Success status. If even one sub-task in a task fails, an exclamation mark with a red circle is displayed next to it, indicating a Failed status.

To view the sub-tasks that are part of a task, click the > on the left of the task to expand the list. You can view the list of sub-tasks with the corresponding status of each sub-task. To collapse the list, click the ▼ on the left of the task.

Hover over a task or sub-task status to see information about the time elapsed and details on the task. For failed tasks, you can view the reason for failure. Failure reasons include errors on hardware, software, interfaces, licenses, and CA certificates.

Table 134: Fields on the Job Status Page

Field	Description
Name	Name of the job.
State	Status of the job - Success or Failed.
Job ID	Unique job ID.
Start Time	Time when the job started.
End Time	Time when the job ended.
Tasks Succeeded	Count of the number of tasks that succeeded in the job. .
Tasks Failed	Count of the number of tasks that failed in the job.

RELATED DOCUMENTATION

[About the Jobs Page | 854](#)

[Add Devices | 149](#)

[View and Manage Device Configuration | 165](#)

Analytics

IN THIS CHAPTER

- [Configure Routers to Send JTI Telemetry Data and RPM Statistics to the Data Collector | 858](#)
- [Configure Routers to Advertise Link Statistics through BGP-LS | 861](#)
- [NetFlow Collector Overview | 864](#)
- [Collect Analytics Data Overview | 873](#)
- [View Analytics Data | 882](#)

Configure Routers to Send JTI Telemetry Data and RPM Statistics to the Data Collector

Junos Telemetry Interface (JTI) sensors generate data (LSP traffic data, logical and physical interface traffic data) from the Packet Forwarding Engine (PFE), and will only send probes through the data plane. So, in addition to connecting the routing engine to the management network, a data port must be connected to the collector on one of your devices. The rest of the devices in the network can use that interface to reach the collector.

NOTE: You must use Junos OS Release 15.1F6 or later for the analytics feature.

To configure the routers, use the following procedure:

1. Configure the devices for telemetry data. On each device, the following configuration is required. The device needs to be set to enhanced-ip mode, which might require a full reboot.

NOTE: Use default remote-port 4000 and set the remote-address to the Paragon Insights virtual IP address.

```

set chassis network-services enhanced-ip
set services analytics streaming-server ns remote-address 192.168.10.100
set services analytics streaming-server ns remote-port 4000
set services analytics export-profile ns local-address 10.0.0.10
set services analytics export-profile ns reporting-rate 2
set services analytics export-profile ns format gpb
set services analytics export-profile ns transport udp
set services analytics sensor ifd server-name ns
set services analytics sensor ifd export-name ns
set services analytics sensor ifd resource /junos/system/linecard/interface/
set services analytics sensor ifl server-name ns
set services analytics sensor ifl export-name ns
set services analytics sensor ifl resource /junos/system/linecard/interface/logical/usage/
set services analytics sensor lsp server-name ns
set services analytics sensor lsp export-name ns
set services analytics sensor lsp resource /junos/services/label-switched-path/usage/
set services analytics sensor sr-te-color server-name ns
set services analytics sensor sr-te-color export-name ns
set services analytics sensor sr-te-color resource /junos/services/segment-routing/traffic-
engineering/ingress/usage/
set services analytics sensor sid server-name ns
set services analytics sensor sid export-name ns
set services analytics sensor sid resource /junos/services/segment-routing/sid/usage/
set services analytics sensor sr-te-tunnels server-name ns
set services analytics sensor sr-te-tunnels export-name ns
set services analytics sensor sr-te-tunnels resource /junos/services/segment-routing/traffic-
engineering/tunnel/ingress/usage/
set protocols mpls sensor-based-stats
set protocols source-packet-routing telemetry statistics

```

In this configuration, the remote address is the IP address of the collector (reachable though a data port). The local address should be the loopback, or router-id, whichever is configured on the device profile to identify the device.

- Bandwidth sizing and container LSPs are supported for SR-TE LSPs. Junos OS release 19.2R1 or later is required for this functionality. There is additional configuration required on the router to enable collection of segment routing data. For example:

```
set groups jvision services analytics sensor sr-te-tunnels server-name ns
set groups jvision services analytics sensor sr-te-tunnels export-name ns
set groups jvision services analytics sensor sr-te-tunnels resource /junos/services/segment-
routing/traffic-engineering/tunnel/ingress/usage/
```

- Real-time performance monitoring (RPM) enables you to monitor network performance in real time and to assess and analyze network efficiency. To achieve this, RPM exchanges a set of probes with other IP hosts in the network for monitoring and network tracking purposes.

You must configure RPM probes to measure the interface delays.

The following example shows the configuration of probes out of interface ge-0/1/1.0 to the remote address 10.101.105.2. This remote address should be the IP address of the node at the other end of the link.

NOTE: The test name must match the interface being measured (test ge-0/1/1.0, in this example).

```
set services rpm probe northstar-ifl test ge-0/1/1.0 target address 10.101.105.2
set services rpm probe northstar-ifl test ge-0/1/1.0 probe-count 11
set services rpm probe northstar-ifl test ge-0/1/1.0 probe-interval 5
set services rpm probe northstar-ifl test ge-0/1/1.0 test-interval 60
set services rpm probe northstar-ifl test ge-0/1/1.0 source-address 10.101.105.1
set services rpm probe northstar-ifl test ge-0/1/1.0 moving-average-size 12
set services rpm probe northstar-ifl test ge-0/1/1.0 traps test-completion
set services rpm probe northstar-ifl test ge-0/1/1.0 hardware-timestamp
set services rpm probe northstar-ifl test ge-0/1/1.0 probe-type icmp-ping-timestamp
```

RELATED DOCUMENTATION

| [About the Tunnel Tab](#) | 710

Configure Routers to Advertise Link Statistics through BGP-LS

You can use Border Gateway Protocol-Link State (BGP-LS) to obtain information about the link delay and link delay variation (variation in the measured delay between consecutive readings) from a network. In the previous releases, the link delay information was obtained from real-time performance monitoring (RPM) probes configured on the routers.

The path computation element (PCE) in Paragon Pathfinder uses the measured link delays to compute the end-to-end LSP delay as a sum of all link delays in an LSP path. If a maximum delay is configured for the LSP and if the computed link delay violates the configured maximum delay, the PCE computes and reroutes the LSP through a path that has a link delay within the configured maximum delay.

NOTE: You can configure Juniper Networks routers to send measured link delay through BGP-LS only if:

- Junos OS 21.3 is running on the router.
- IS-IS protocol is configured on a point-to-point link

To configure a Juniper Networks routers to send the measured delay and delay variation on a link through BGP-LS:

1. Enable Two-Way Active Management Protocol (TWAMP) on the routers (Router A and Router z) at the ends of the link:

```
set services rpm twamp server authentication-mode none
set services rpm twamp server light
```

2. On each router, configure the interfaces on which you want to measure the delay. For example:

```
set protocols isis interface ge-0/1/1.0 delay-measurement
set protocols isis interface ge-0/1/1.0 point-to-point
```

3. If RPM probes are enabled on the routers, disable the RPM probes to ensure that the routers do not send link delay statistics obtained through RPM probes:

```
deactivate services rpm probe northstar-if1
```

NOTE: Refer to ["Configure Routers to Send JTI Telemetry Data and RPM Statistics to the Data Collector"](#) on page 858 for the RPM probes that might be configured on the router.

4. (Optional) View the measured delay in the traffic engineering database (TED) database and containerized routing protocol daemon (cRPD) by using the following commands:

- To view the measured delay in the TED database of a router, execute the following command in the router:

```
show ted database extensive ip-address-of -router
```

The output lists the average delay, minimum delay, maximum delay, and delay variation in milliseconds (ms) as follows:

```
...
Local interface index: 368, Remote interface index: 362
Color: 0x6 red blue
Metric: 10
IGP metric: 10
Average delay: 1138
Minimum delay: 643
Maximum delay: 4401
Delay variation: 1565
Static BW: 10Mbps
...
```

- To view the measured delay received by cRPD, execute the following command in cRPD:

```
show route table lsdist.0 hidden extensive
```

The output lists the average delay, minimum delay, maximum delay, and delay variation as follows:

```
...
Metric: 10
TE Metric: 10
Average delay: 1094
Minimum delay: 700
Maximum delay: 3487
Delay variation: 1396
```

```


SRLG membership:
Unknown-0x64
Label: 31, Flags: 0x30, Weight: 0
Localpref: 100
...

```

You can view the link delay statistics in the Paragon Automation GUI as follows:

- Right-click a link label in the network topology (**Network > Topology**) and select **Measured Delay A::Z**. The measured delay is listed on the links as shown in [Figure 69 on page 863](#).

Figure 69: Measured Link Delay Indicated on Links in the Network Topology



- The Measured Delay A and Measured Delay Z columns in the Links tab of the Network Information table display the measured delay on respective ends of the link; see [Figure 70 on page 864](#).

Figure 70: Measured Delay Displayed in the Network Information Table

Link	Node A	Node Z	Status	Interface A	Interface Z	Bandwidth A	Bandwidth Z	Measured Delay A	Measured Delay Z
<input checked="" type="checkbox"/>	vmx102	vmx105	Up	ge-0/1/2.0	ge-0/1/2.0	10M	10M	1.227	0.925
<input checked="" type="checkbox"/>	vmx102	vmx106	Up	ge-0/1/3.0	ge-0/1/3.0	10M	10M	1.13	2.317

RELATED DOCUMENTATION

[View Analytics Data](#) | 882

NetFlow Collector Overview

IN THIS SECTION

- [Demand Generation](#) | 865
- [NetFlow Collector Requirements](#) | 866

NetFlow collector is a data collection tool in Paragon Pathfinder.

NetFlow collector uses the netflowd microservice in the northstar namespace to collect and report data about traffic flow in the network. Netflowd is automatically installed as part of the Pathfinder installation package. Netflowd receives the NetFlow data from the routers, decodes the records, and aggregates the data. Netflowd uses this aggregated data to create demands, which indicate the amount of traffic flow in the network. Netflowd stores the data in the Time Series Database (TSDB) and shares the data with the Path Computation Server (PCS). The aggregated data is used to generate Demand reports that are available in Paragon Pathfinder (**Reports > Demand**). These reports provide information on network traffic. This data is also used in Paragon Planner to generate Demand reports, plan, and model the network.

NOTE:

- The PCS monitors traffic from the autonomous systems (AS) and VPNs.

- The PCS supports both IPv4 and IPv6 traffic.

Pathfinder leverages the Junos OS implementation of flow monitoring and aggregation by using NetFlow Version 9 and Version 10 (IPFIX) flow templates. See the following Junos OS documentation for background:

- *Configuring Flow Aggregation to Use Version 9 Flow Templates.*
- *Configuring Flow Aggregation to Use IPFIX Flow Templates on MX, vMX and T Series Routers, EX Series Switches and NFX250.*
- *Configuring Flow Aggregation to Use IPFIX Flow Templates on PTX Series Routers.*

Demand Generation

Netflowd uses four aggregation keys (obtained from the NetFlow data that netflowd collects) to generate the demands:

- Ingress provider edge (PE) device (that is the device reporting the flow)
- BGP next hop IP address
- Routing table name
 - When this key is present, it is the name of the VRF for which the ingress interface is configured.
 - This key is absent if no VPN is associated with the demand. In this case, the ingress interface is configured in the default routing table.
 - This key is displayed as **NONE** if netflowd is not able to determine whether the ingress interface is configured in the default routing table or on a VRF. That would happen, for example, if the PCS was not able to collect the snmp-indexes for the interfaces.
- Specification of IPv4 (displayed as IP in the Demand tab of the network information table) or IPv6

The values of the keys are indicated in the names of the demands which are displayed in the Name column of the Demand tab in the network information table. Here are some examples:

- vmx102_10.1.0.10/32_vpn100_IP
- vmx102_10.1.0.10/32_IP (if no VPN is associated with the demand)
- vmx102_10.1.0.10/32_NONE_IP (if it is unknown whether the ingress interface is configured on the default routing table or on a VRF)

NetFlow Collector Requirements

To use NetFlow collector in Pathfinder, you must:

- Install Nginx Ingress Controller when you install the Infrastructure component. See [Install Multi-Node Cluster on CentOS](#) and [Install Multi-Node Cluster on Ubuntu](#).
- Configure the network routers for flow monitoring (NetFlow v9 or v10). See "[Configuration on the Network Routers](#)" on page 866.
- Run the device collection task periodically to create and maintain an accurate VPN model in Pathfinder. We recommend that you run the device collection task at least once daily. See "[Add a Device Collection Task](#)" on page 978.

(Optional) Customize NetFlowd parameters from the CLI. See "[Customize Netflowd Parameters from the CLI](#)" on page 870.

Configuration on the Network Routers

To use NetFlow collector in Pathfinder, you must configure the network routers for flow monitoring (NetFlow v9 or v10) according to the router's operating system documentation.

NOTE: Currently, only Juniper Networks devices and Cisco IOS-XR devices can be configured with NetFlow v9 and v10.

Here are some important considerations to keep in mind when you configure the routers:

- The NetFlow process (netflowd) identifies the device, which is reporting the flow, through the source address (`inline-jflow` statement). Configure this parameter as the router's loopback address.
- The flow-active-timeout parameter has a default value of 60 seconds. We recommend keeping it at 60 seconds or less.
- Configure the flow-server's IP address as the virtual IP (VIP) address that is configured for the Nginx Ingress Controller.

The following Junos OS example shows the NetFlow v9 configuration statements.

At the interface hierarchy level:

```
interfaces {
  ge-0/0/1 {
    unit 0 {
      family inet {
```

```

        sampling {
            input;
        }
        address 10.0.21.1/24;
    }
}
}
}
}

```

At the forwarding-options hierarchy level:

```

forwarding-options {
  sampling {
    instance {
      nfv9-ipv4 {
        input {
          rate 1;
          run-length 0;
        }
        family inet {
          output {
            flow-inactive-timeout 15;
            flow-active-timeout 60;
            flow-server 172.16.18.1 {
              port 9000;
              version9 {
                template {
                  nfv9-ipv4;
                }
              }
            }
            inline-jflow {
              source-address 10.1.0.104;
            }
          }
        }
      }
    }
  }
}

```


At the chassis hierarchy level:

```
chassis {
  network-services enhanced-ip;
  fpc 0 {
    sampling-instance nfv9-ipv4;
  }
}
```

At the services hierarchy level:

```
services {
  flow-monitoring {
    version9 {
      template nfv9-ipv4 {
        nexthop-learning {
          enable;
        }
        template-refresh-rate seconds 60;
        option-refresh-rate seconds 60;
        ipv4-template;
      }
    }
  }
}
```

The following Junos OS example shows NetFlow v10 configuration statements.

At the interface hierarchy level:

```
interfaces {
  ge-0/0/1 {
    unit 0 {
      family inet {
        sampling {
          input;
        }
        address 10.0.21.1/24;
      }
    }
  }
}
```

```

    }
}

```

At the forwarding-options hierarchy level:

```

forwarding-options {
  sampling {
    instance {
      nfv10-ipv4 {
        input {
          rate 1;
          run-length 0;
        }
        family inet {
          output {
            flow-inactive-timeout 15;
            flow-active-timeout 60;
            flow-server 172.16.18.1 {
              port 9000;
              version-ipfix {
                template {
                  nfv10-ipv4;
                }
              }
            }
            inline-jflow {
              source-address 10.1.0.104;
            }
          }
        }
      }
    }
  }
}

```

At the chassis hierarchy level:

```

chassis {
  network-services enhanced-ip;
  fpc 0 {
    sampling-instance nfv10-ipv4;
  }
}

```

```

    }
}

```

At the services hierarchy level:

```

services {
  flow-monitoring {
    version-ipfix {
      template nfv10-ipv4 {
        nexthop-learning {
          enable;
        }
        template-refresh-rate {
          seconds 60;
        }
        option-refresh-rate {
          seconds 60;
        }
        ipv4-template;
      }
    }
  }
}

```

Customize Netflowd Parameters from the CLI

The parameters related to netflowd are configured by default and you can view these parameters from the CLI. Optionally, you can customize these parameters as well. [Table 135 on page 870](#) describes the netflowd parameters that you can customize from the CLI.

Table 135: Netflowd Parameters

Parameter	Command	Command
enable-ssl	set northstar analytics netflowd enable-ssl	Configure this parameter to enable netflowd to establish a Secure Socket Layer (SSL) connection to the native datastore.

Table 135: Netflowd Parameters (Continued)

Parameter	Command	Command
logging-parameters	set northstar analytics netflowd logging-parameters	Configure the level of information that is captured in the log file. The default level is info . If you want more information to be included in the log file, you can set the level to debug . The log file will include all the flows which are received from each device, identified by the source IP address. You can also view, for each flow, all the fields that netflowd processes and parses.
default-sampling-interval	set northstar analytics netflowd default-sampling-interval	Configure the default sampling interval that is used if the router does not provide the interval in the Template FlowSet. Default: 1.
publish-interval	set northstar analytics netflowd publish-interval	Configure the interval (in seconds or minutes) to publish records to both the TSDB and the PCS. Traffic is aggregated per publishing interval. This value must be equal to or greater than the reporting time configured in the router (flow-active-timeout value) to ensure that for every publishing interval, all active flows are reported. Default: 60s.
notify-final-bandwidth-on-inactive-flow	set northstar analytics netflowd notify-final-bandwidth-on-inactive-flow	Configure this parameter to enable netflowd to send one final update after a flow is no longer active, reporting the bandwidth as 0. By default, this parameter is not configured. So, the bandwidth value is not reported once a flow becomes inactive; the last reported active value is the last value displayed.
aggregate-by-prefix	set northstar analytics netflowd aggregate-by-prefix	Configure this parameter to enable netflowd to aggregate all traffic from a specific ingress provider edge (PE) router to a specific destination (prefix) within the specified period. By default, NetFlow aggregates traffic by PE routers, but for some applications (such as Egress Peer Engineering and Ingress Peer Engineering), you would want the traffic to be aggregated by prefix.

Table 135: Netflowd Parameters (Continued)

Parameter	Command	Command
stats-interval	set northstar analytics netflowd stats-interval	Configure the interval (in seconds) at which statistics are printed to the log file. By default, the interval is not configured, so the statistics are not printed to the log file.
generate-as-demands	set northstar analytics netflowd generate-as-demands	Configure this parameter to enable netflowd to generate AS demands. By default, this parameter is not configured. So, AS demands are not displayed through REST APIs or in Demand reports in the GUI, even if valid NetFlow records are being exported.
top-prefixes	set northstar analytics netflowd top-prefixes	Configure the number of prefixes (in terms of the aggregated traffic volume) to be exported. Range: 1 through 10,000
top-prefixes-export-ticks	set northstar analytics netflowd top-prefixes-export-ticks	Configure the number of intervals above which traffic is aggregated for the top N prefixes, where the export interval length is determined by the publish-interval parameter. Example: If you set the publish-interval as 60s and top-prefixes-export-ticks as 5, the top N prefixes are exported (published) every 5 minutes (5x60s = 5m).
workers	set northstar analytics netflowd workers	Configure the number of processes to be started. When set to 0, it takes the value of the number of cores in the system. Default: 1

RELATED DOCUMENTATION

| [Add a Demand Reports Task](#) | 985

Collect Analytics Data Overview

IN THIS SECTION

- [Workflow to Collect Device Statistics | 877](#)
- [Sample Configuration for collecting Telemetry from Cisco IOS XR Devices | 880](#)
- [Data Summarization and Retention Policy | 881](#)

Paragon Automation supports collecting analytics data from Juniper Networks, Nokia, and Cisco IOS XR devices.

- [Table 136 on page 873](#) lists the data collected from Juniper Networks devices.
- [Table 137 on page 875](#) lists the data collected from Cisco IOS XR devices.
- [Table 138 on page 877](#) lists the data collected for Nokia devices.

Table 136: Data Collected for Analytics: Juniper Devices

Data	Collected by Default	Ingest Type	Rules Used to Collect Data
Label-switched Path (LSP) statistics [RSVP, Segment Routing (both colored and non-colored)]* NOTE: Currently, Paragon Pathfinder does not support the collection of SRv6 LSP statistics.	Yes	JTI	controller-telemetry /ctrl-label-switched-path controller-telemetry /ctrl-label-switched-path-aggregation

Table 136: Data Collected for Analytics: Juniper Devices (Continued)

Data	Collected by Default	Ingest Type	Rules Used to Collect Data
Logical and physical interface statistics	Yes	JTI	controller-telemetry /ctrl-logical-interface controller-telemetry /ctrl-logical-interface-aggregation controller-telemetry /ctrl-physical-interface controller-telemetry /ctrl-physical-interface-aggregation
Segment ID (SID)	Yes	JTI	controller-telemetry /ctrl-sr-sid
Link Latency	Yes	iAgent	controller-telemetry /ctrl-link-latency
LSP Latency	Yes	Derived from Syslog	controller-telemetry /ctrl-lsp-latency
Quality of Service (QoS)	No	JTI	controller-telemetry /ctrl-egress-queue-interface controller-telemetry /ctrl-ingress-queue-interface
Label Distribution Protocols (LDP) demands	No	iAgent	controller-telemetry /ctrl-ldp-demand-stats
Netflow	Yes	Not Applicable	controller.telemetry/ctrl-demand-stats controller.telemetry/ctrl-as-demand-stats

Table 136: Data Collected for Analytics: Juniper Devices (Continued)

Data	Collected by Default	Ingest Type	Rules Used to Collect Data
------	----------------------	-------------	----------------------------

NOTE:

- For collecting SR LSP statistics from a Juniper Networks device, the device should be running Junos OS Release 20.2 or a later version.
- For collecting LSP latency, the `lsp-latency-interval` must be set. The default value is 180s. You can modify this value; see ["Modify Pathfinder Settings From the GUI" on page 206](#) or ["Modify Pathfinder Settings From the Pathfinder CLI" on page 198](#) for details.
- The segment routing LSPs created by using node SIDs might have multiple paths. Therefore, delay for such LSPs is not deterministic. Thus, Paragon Pathfinder cannot calculate and display the LSP delay.

NOTE: For Juniper devices, the Path Computation Element Protocol (PCEP) reports the route metric. For more information, see [Basic LSP Configuration](#). You can set the route metric for an LSP and configure the global preference of Paragon Pathfinder from the **Configuration > Network Settings > Pathfinder Settings > Path Computation Server** page.

Table 137: Data Collected for Analytics: Cisco Devices

Data	Collected by Default	Ingest Type	Rules Used to Collect Data
Label-switched Path (LSP) statistics [RSVP, segment routing]	Yes	SNMP The Management Information Base (MIB) used is MPLS-TE-STD-MIB::mplsTunnelTable).	controller-telemetry /ctrl-label-switched-path-snmp-aggregation
Label-switched Path (LSP) statistics [segment routing, color]	No	SNMP The Management Information Base (MIB) used is IF-MIB::ifXTable.	controller-telemetry /ctrl-label-switched-path-sr-color-snmp-cisco-aggregation

Table 137: Data Collected for Analytics: Cisco Devices (Continued)

Data	Collected by Default	Ingest Type	Rules Used to Collect Data
Logical and physical interface statistics	Yes	<ul style="list-style-type: none"> SNMP (For devices running IOS XR Release 7.1.0 and prior) The MIB used is IF-MIB::ifXTable. OpenConfig (For devices running IOS XR Release 7.1.1 and later) 	<ul style="list-style-type: none"> SNMP: <ul style="list-style-type: none"> controller-telemetry /ctrl-logical-interface-snmp-aggregation controller-telemetry /ctrl-physical-interface-snmp-aggregation OpenConfig: <ul style="list-style-type: none"> controller-telemetry /ctrl-logical-interface controller-telemetry /ctrl-logical-interface-aggregation controller-telemetry /ctrl-physical-interface controller-telemetry /ctrl-physical-interface-aggregation
Netflow	Yes	Not Applicable	<ul style="list-style-type: none"> controller.telemetry/ctrl-demand-stats controller.telemetry/ctrl-as-demand-stats

NOTE: Cisco devices do not send color attribute over PCEP. You have to manually configure the color attribute for SR colored LSPs or run the device collection on the device.

Sample configuration:

API - `https://{UI-IP}/traffic-engineering/api/topology/v2/1/te-lsps/{LSP-INDEX}`

Operation - PATCH

Body - `[{"op": "replace", "path": "/plannedProperties/color", "value": <color>}]`

For Cisco IOS XR devices, the reported metric is the sum of the IGP metrics on all outgoing interfaces along a particular path from the source to the destination.

Table 138: Data Collected for Analytics: Nokia Devices

Data	Collected by Default	Ingest Type	Rules Used to Collect Data
Label-switched Path (LSP) statistics [RSVP, segment routing]	Yes	SNMP The Management Information Base (MIB) used is MPLS-TE-STD-MIB::mplsTunnelTable).	controller-telemetry /ctrl-label-switched-path-snmp-nokia-aggregation
Logical and physical interface statistics	Yes	SNMP The Management Information Base (MIB) used is IF-MIB::ifXTable.	SNMP: controller-telemetry /ctrl-logical-interface-snmp-aggregation controller-telemetry /ctrl-physical-interface-snmp-aggregation

NOTE: The default SNMP collection interval is 300 seconds. You can modify the interval from the `ctrl-snmp-frequency` field in the Frequency profile tab on the Ingest Settings page (**Configuration > Data Ingest > Settings**). See ["Manage Frequency Profiles"](#) on page 469 for details.

For Nokia devices, the reported metric is the sum of the IGP metrics on all outgoing interfaces along a particular path from the source to the destination.

Workflow to Collect Device Statistics

Paragon Automation contains a default *controller* device group and a default *controller* playbook. Devices must be added to the controller device group so that analytics data can be collected from these devices. The controller playbook defines the rules for collecting the analytics data.

The following is the workflow to collect analytics data in Paragon Automation:

1. If not added already, add the devices from which you want to collect data, to the Controller device group. See ["Edit a Device Group"](#) on page 183.

2. Run the device collection task to correlate the links to the interfaces displayed in the topology view (Network > topology); see ["Add a Device Collection Task" on page 978](#).
3. For a Juniper Networks device, configure the devices to send JTI telemetry data and RPM statistics; see ["Configure Routers to Send JTI Telemetry Data and RPM Statistics to the Data Collector" on page 858](#).
4. For a Cisco IOS XR device, configure OpenConfig or SNMP for collecting interface statistics:
 - a. Access the CLI of the Cisco IOS XR device.
 - b. Do one of the following:
 - To configure OpenConfig on Cisco IOS XR devices, add the following configuration:

```

grpc
port 32767
no-tls
!
telemetry model-driven
sensor-group JTIMON_INTERFACE
sensor-path openconfig-interfaces:interfaces
!
subscription hbot_interfaces_
sensor-group-id JTIMON_INTERFACE sample-interval 10000
!
subscription hbot_interfaces_interface_state_
sensor-group-id JTIMON_INTERFACE sample-interval 10000
!
subscription hbot_interfaces_interface_subinterfaces_subinterface_state_
sensor-group-id JTIMON_INTERFACE sample-interval 10000
!
!

```

- To configure SNMPv2 on Cisco IOS XR devices, add the following configuration:

```
snmp-server community public R0
```

- To configure SNMPv3 on Cisco IOS XR devices, add the following configuration:

NOTE: The following is a sample configuration.

```
snmp-server view ViewDefault iso included
snmp-server group GrpMonitoring v3 priv read ViewDefault
snmp-server user UserJustMe GrpMonitoring v3 auth sha AuthPass1 priv aes 128 PrivPass2
```

After you add the SNMPv3 configuration, configure the SNMPv3 parameters from the Devices page (**Configuration > Devices**). See ["Configure SNMP Ingest" on page 499](#) for details.

- c. Commit the configuration.

```
commit
```

5. Add optional rules to the controller playbook in Paragon Automation.

The rules for collecting LSP statistics, interface statistics, SID statistics, and link latency are added to the controller playbook by default.

If you want to collect QoS and LDP demand statistics, you must add the rules for collecting these statistics manually. To add rules to the controller playbook.

- a. Click the **Configuration > Playbooks** icon in the left navigation menu.

The Playbooks page appears.

- b. Click the **Controller** playbook.

The Edit Playbook Controller page appears.

- c. In the **Rules** field, click and select the following rules:

- controller-telemetry/ctrl-egress-queue-interface and controller-telemetry/ctrl-ingress-queue-interface for collecting QoS data.
- controller-telemetry/ctrl-ldp-demand-stats for collecting LDP demand statistics.

- d. Click **Save and Deploy**.

A confirmation message appears that the playbook is successfully saved and deployed on the devices in the controller device group.

For more details on editing a playbook, see ["Edit a Playbook" on page 310](#)

- e. Pause the current playbook instance assigned to the controller device group and add the new playbook instance. see ["Manage Playbook Instances" on page 312](#) for details.
- f. (Optional) Delete the paused playbook instance.

6. Commit the configuration.

```
commit
```

7. Add the following data rollup summarization profiles to the controller device group, if you added rules for QoS data in step ["5" on page 879](#):

ctrl-egress-queue-interface-rollup-profile and ctrl-ingress-queue-interface-rollup-profile
 controller-default-rollup-profile is added to the controller device group by default.

For details about data summarization profiles, see ["Data Summarization Overview" on page 647](#) and ["Apply Data Summarization Profiles" on page 655](#).

8. Add devices to the collector device group to enable collecting the device data.

Do one of the following to add a device to a device group:

- Edit the controller device group to add more devices; see ["Edit a Device Group" on page 183](#)
- Assign the device to the controller device group; see ["Edit Devices" on page 168](#)

NOTE: The Junos telemetry interface (JTI) sensors use the system ID parameter of Juniper Networks devices to collect data from the devices. So, if not already configured in Paragon Automation, configure the system ID for the Juniper Networks devices. The system ID should be in the following format: `<host_name><jti_ip_address>` and should have the same value as configured on the device; see ["Edit Devices" on page 168](#) for configuring the system ID of a device in Paragon Automation.

Sample Configuration for collecting Telemetry from Cisco IOS XR Devices

The subscription name on a Cisco device must be derived from the *sensor-name* of the Paragon Insights rule. The subscription depends on the *sensor-name* you want to execute on the Cisco device. For example, if the *sensor-name* is `/interfaces-test`, the subscription name needs to be configured in IOS-XR as `hbot_interfaces_test_`.

When Paragon Automation sends the subscription request to the Cisco device:

- A prefix `hbot_` is added to the *sensor-name*, and

- all the non-letter and non-digit characters are replaced with an underscore (_).

NOTE: The conversion will take place only if the *sensor-name* is configured with a leading slash (/).

The following is sample configuration of the Paragon Insights rule for obtaining telemetry from Cisco IOS XR devices:

```
set healthbot topic controller.telemetry rule ctrl-label-switched-path sensor sr-lsp-mdt open-
config sensor-name /jtimon_sr_te
set healthbot topic controller.telemetry rule ctrl-label-switched-path sensor sr-lsp-mdt open-
config frequency 60s
set healthbot topic controller.telemetry rule ctrl-label-switched-path field lsp-stats-bytes
sensor sr-lsp-mdt path /xtc/policy-forwardings/policy-forwarding/stats/bytes
set healthbot topic controller.telemetry rule ctrl-label-switched-path field lsp-stats-counter
sensor sr-lsp-mdt path "/xtc/policy-forwardings/policy-forwarding/@name"
set healthbot topic controller.telemetry rule ctrl-label-switched-path field lsp-stats-packets
sensor sr-lsp-mdt path /xtc/policy-forwardings/policy-forwarding/stats/packets
set healthbot topic controller.telemetry rule ctrl-label-switched-path field to-ip sensor sr-lsp-
mdt path /xtc/policy-forwardings/policy-forwarding/active-lsp/candidate-path/name
```

Since the *sensor-name* is /jtimon_sr_te, the Cisco IOS configuration will be:

```
subscription hbot_jtimon_sr_te_
sensor-group-id sr-te sample-interval 10000
!
```

Data Summarization and Retention Policy

Paragon Automation uses summarization (rollup) profiles to rollup data so that the data is stored in a smaller disk space and performance of the time-series database (TSDB) is improved. Raw network and device data are grouped together into hourly and daily data by using aggregate functions. For more information on data summarization, see ["Data Summarization Overview" on page 647](#).

By default, the data is retained for the following duration:

- Daily rollup data is stored for 1000 days.
- Hourly rollup data is stored for 180 days.

- Raw data is stored for 14 days.

You can modify the default duration for which the rollup data is retained from the Ingest Settings page **Administration > Ingest Settings > Retention Policy**.

The network archive task uses the rolled-up data for providing network statistics. Also, statistics are displayed on the topology graphs by fetching long range data from rolled-up measurements based on the `rollup-query-cutoff-interval` Paragon Pathfinder setting.

`rollup-query-cutoff-interval` defines the time period for which raw data or rolled-up data is fetched. If data is queried for a time period less than or equal to the value set in `rollup-query-cutoff-interval`, raw data is fetched. If data is queried for a time greater than `rollup-query-cutoff-interval`, hourly rolled-up data is fetched. The default value is set to 7 days; that is, by default, rolled-up data is fetched when queried for data greater than 7 days. For information about setting the `rollup-query-cutoff-interval` parameter, see ["Modify Pathfinder Settings From the GUI" on page 206](#) or ["Modify Pathfinder Settings From the Pathfinder CLI" on page 198](#).

View Analytics Data

IN THIS SECTION

- [Analytics Widgets View | 883](#)
- [Network Performance Data | 887](#)
- [Traffic through a Node | 891](#)
- [Traffic and Delay Information for a Link | 891](#)
- [Traffic and Delay Information for a Tunnel | 893](#)
- [Traffic and Delay Information for an Interface | 895](#)
- [Analytics Information on the Topology Map | 896](#)

You can view the analytics data collected for your network so that the data can be interpreted and acted upon. You can view the network analytics data in:

- Widgets on the dashboard; see ["Analytics Widgets View" on page 883](#)
- Performance data in the topology view; see ["Network Performance Data" on page 887](#)
- Network information table

- Node traffic; see ["Traffic through a Node" on page 891](#)
- Link data; see ["Traffic and Delay Information for a Link" on page 891](#)
- Tunnel data; see ["Traffic and Delay Information for a Tunnel" on page 893](#)
- Interface data; see ["Traffic and Delay Information for an Interface" on page 895](#)
- Topology view; see ["Analytics Information on the Topology Map" on page 896](#).

Analytics data is collected in two forms—Historical data and Live aggregated data.

The historical data is used for generating network reports. The live aggregated data for one minute is displayed in the topology view.

Analytics Widgets View

The following widgets on the dashboard display the analytics data:

- Top LSP Sources—This widget lists the devices in a network that have the highest number of LSPs originating from them, and the number of originating LSPs, as shown in [Figure 71 on page 883](#). All the devices in the network from which data is collected are listed in the decreasing order of the number of originating LSPs.


Figure 71: Top LSP Sources

Top LSP Sources		
Index	Device	Count
1	ios-xr8	10
2	vmx101	4
3	vmx104	3
4	vmx102	3
5	vmx103	3
6	ios-xr9	3

Last updated: Jul 2, 2021, 10:54:44 AM ?


- **Top LSP Destinations**—This widget lists the top 10 devices in a network that have the highest number of LSPs terminating on them and the number of terminating LSPs. All the devices in the network from which data is collected are listed in the decreasing order of the number of terminating LSPs.
- **LSP Summary**—This widget displays a bar graph that indicates the number of primary, secondary, and standby LSPs that are Up (or Active) and Down in a network as shown in [Figure 72 on page 884](#).

Figure 72: LSP Summary




- **Top 10 Traffic-Devices**—This widget displays a bar graph of the top 10 devices through which maximum traffic (in bps) is flowing, as shown in [Figure 73 on page 885](#).

Figure 73: Top 10 Devices with Maximum Traffic




- Top 10 Traffic-Interfaces—This widget displays a bar graph of the top 10 interfaces through which the maximum traffic (in bps) is flowing, as shown in [Figure 74 on page 885](#).

Figure 74: Top 10 Interfaces with Maximum Traffic




- Top 10 Traffic-LSPs—This widget displays a bar graph of the top 10 LSPs through which maximum traffic ((in bps) is flowing, as shown in [Figure 75 on page 886](#).

Figure 75: Top 10 LSPs with Maximum Traffic



- Top 10 Delay-Interfaces—This widget displays a bar graph of the top 10 interfaces with the maximum measured delay.
- LSP Hop Count—This widget displays a bar graph of the hop count for each type of LSP (PCE-initiated, Delegated, and Device-controlled), as shown in [Figure 76 on page 887](#).

Figure 76: LSP Hop Count




Network Performance Data

You can view the network performance in the topology view (**Network > Topology**). The topology view is updated once every minute. You can view the following network performance data on the topology view:

- **RSVP Live Utilization**—Shows the bandwidth utilization as reported by the devices from the traffic engineering database (through IGP).
- **RSVP Utilization**—Shows the color-coded links based on the reserved bandwidth, as calculated by Paragon Pathfinder from the existing reservations (reported by using PCEP, NETCONF, or both). The color legend is displayed at the bottom right corner of the topology view.
- **Interface Utilization**—Shows the last measured interface traffic pushed from the devices to the data collectors. The interface utilization is displayed as the percentage utilization of the links in the format percentage A-Z::percentage Z-A, as shown in [Figure 77 on page 888](#).

Figure 77: Topology View Displaying Interface Utilization




The presence of two different colors on a link indicates that the utilization in one direction (A to Z) is different from the utilization in the other direction (Z to A). The half of the link originating from a certain node is colored according to the link utilization in the direction from that node to the other node.

- Calculated Interface Delay -Shows the interface delay data calculated by the path computation server (PCS) based on the distance between the nodes.
- Measured Interface Delay—Shows the delay introduced by an interface, based on inputs from RPM probes.
- Node Ingress Traffic—Shows the bandwidth consumed by the ingress traffic of nodes in the network. The nodes are color-coded based on the amount of traffic flowing through them, as shown in [Figure 78 on page 889](#). For example, nodes vmx107 and vmx105 are colored red because the ingress traffic to these nodes consume more than 80% of the node capacity. A legend, present at the right-corner of the topology view, explains the color code.


The links between the nodes indicate the percentage of bandwidth consumed by ingress traffic of each node, in the link, in the format percentage A-Z :: percentage Z-A, as shown in [Figure 78 on page 889](#). For example, on the link between vmx103 and vmx107, 74.93% of the bandwidth is consumed by the ingress traffic to vmx103, and 20.45% of the bandwidth is consumed by the ingress traffic to vmx107.

Figure 78: Node Ingress Traffic



- Node Egress Traffic—Displays the bandwidth consumed by egress traffic of each node, in the link, by the egress traffic of nodes in the network. The bandwidth consumed is displayed in the format percentage A-Z :: percentage Z-A. The nodes are color-coded based on the legend present at the right-corner of the topology view.
- Interface Delay—Displays the delay introduced by the interface in milliseconds (ms), as measured by the PCS based on inputs from RPM probes configured on the interfaces. [Figure 79 on page 890](#) shows a topology view with the measured interface delay displayed on the links in the format delay A-Z :: delay Z-A.

Figure 79: Interface Delay



- Historical Interface Utilization—Shows the average interface utilization over a six day period; see [Figure 77 on page 888](#).

NOTE:

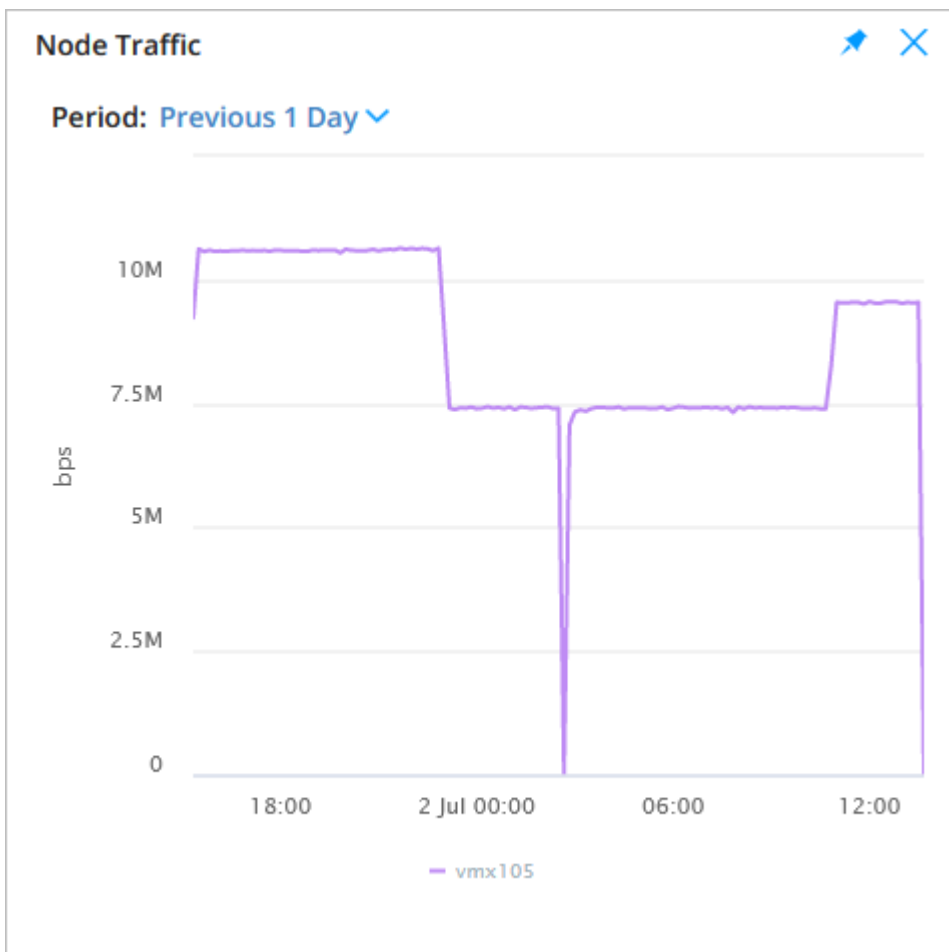
- Node egress traffic, node ingress traffic, and interface delay are historical data. You can view the current interface utilization and also historical interface utilization.
- The RSVP and live RSVP utilization data displayed might be different due to the way the live utilization is reported. To reduce flooding, devices only report reservation changes whenever some thresholds are crossed. In addition, the RSVP utilization can take into account other sources of information that use bandwidth, but don't reserve it, such as the case when using Source Packet Routing (SPRING) LSPs, or LSPs that are configured and placed but are still in the process of being pushed to the network.

Traffic through a Node

For a node, you can view the data about egress traffic through the node in the network information table.

To view egress traffic through a node, in the Nodes tab of the network information table, select the node for which you want to view the traffic and click **View > Node Traffic**. A graph of the traffic (in bps) at different times is displayed, as shown in [Figure 80 on page 891](#).

Figure 80: Egress Traffic through a Node



You can view the graph for the past 3 hours, past 1 day, past 1 week, or define your own custom time period in the *Custom* option. You can also pin the graph.

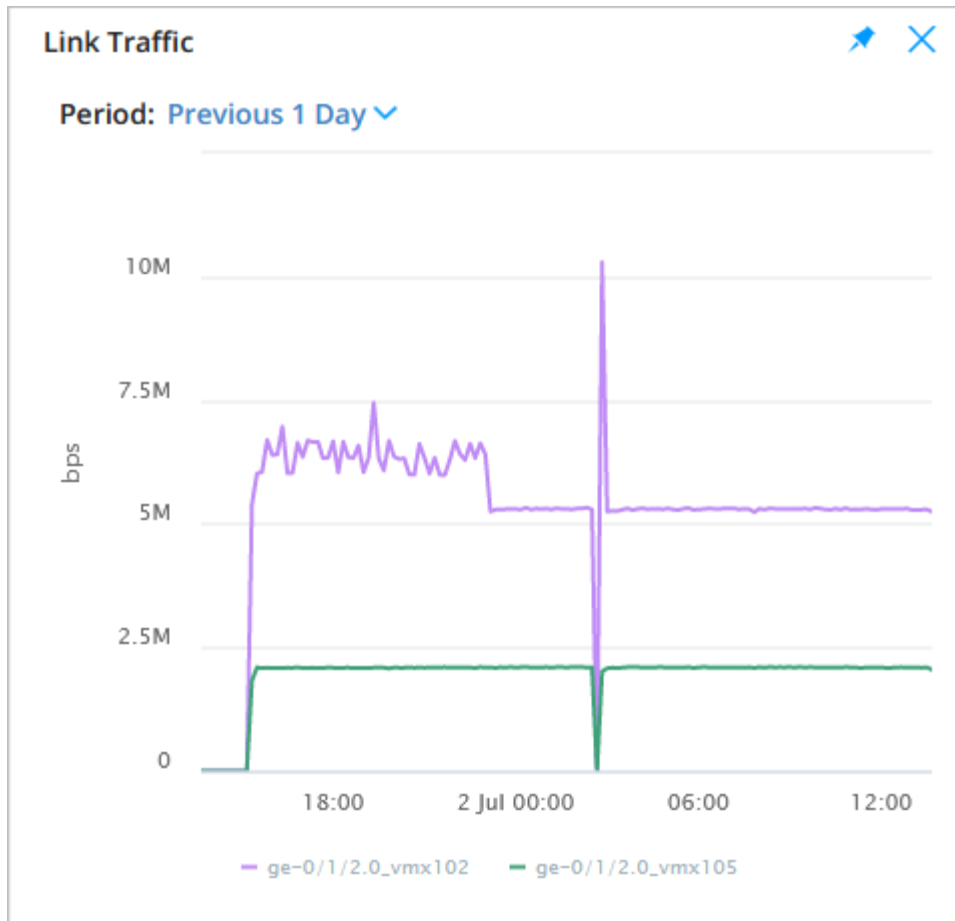
Traffic and Delay Information for a Link

For links, you can view the link traffic and link delay data in the network information table.

To view the traffic and delay data collected for a link, in the Link tab of the network information table, select the link and click:

- **View > Link Traffic** to view the traffic through the link in both the nodes of the link, as shown in [Figure 81 on page 892](#).


Figure 81: Link Traffic



You can view the traffic data for the past 3 hours, past 1 day, past 1 week, or define your own custom time period in the *Custom* option. You can also pin the graph.

- **View > Link Delay** to view the delay introduced in the link at different times, as shown in [Figure 82 on page 893](#).

Figure 82: Link Delay



You can view the link delay data for the past 3 hours, past 1 day, past 1 week, or define your own custom time period in the *Custom* option. You can also pin the graph.


Traffic and Delay Information for a Tunnel

For tunnels, you can view data related to the traffic through the tunnel and delay introduced in the tunnel.

To view the traffic and delay data collected for a tunnel, in the Tunnels tab of the network information table, select the tunnel for which you want to view the data and click::

- **View > LSP Traffic** to view the traffic through the tunnel, as shown in [Figure 83 on page 894](#).

Figure 83: Tunnel Traffic



NOTE: You can view LSP traffic data only for Juniper Networks devices.

You can view the link delay data for the past 3 hours, past 1 day, past 1 week, or define your own custom time period in the *Custom* option. You can also pin the graph.

- **View > Delay** to view the delay introduced in the tunnel, as calculated by the PCS based on inputs from RPM probes, configured on the interfaces.

At any given time, the path computation element (PCE) is aware of the paths of all LSPs in the network. Periodically, the controller uses the reported link delays to compute the end-to-end LSP delay as the simple sum of all link delays in the LSP path.

NOTE: The `lsp-latency-interval` parameter must be set in the Pathfinder configuration to view the delay information for a tunnel. To set the `lsp-latency-interval` parameter, see ["Modify Pathfinder Settings From the Pathfinder CLI"](#) on page 198 or ["Modify Pathfinder Settings From the GUI"](#) on page 206.


Traffic and Delay Information for an Interface

For interfaces, you can view data related to the traffic through the interface and the delay introduced in the interface, in the network information table.

To view the traffic and delay data for an interface, in the Interface tab of the network information table, select the interface for which you want to view the data and click:

- **View > Traffic** to view the traffic through the interface, as shown in [Figure 84 on page 895](#).


Figure 84: Interface Traffic



You can view the interface traffic data for the past 3 hours, past 1 day, past 1 week, or define your own custom time period in the *Custom* option. You can also pin the graph.

- **View > Delay** to view the delay introduced in the interface, as shown in [Figure 85 on page 896](#).

Figure 85: Interface Delay



You can view the interface delay data for the past 3 hours, past 1 day, past 1 week, or define your own custom time period in the *Custom* option. You can also pin the graph.

Interface delay information is available only if:

- RPM probes are configured on the nodes.
- The `rpm-log.slax` script is loaded, to send the results of the probes to the data collectors.

NOTE: Paragon Automation does not automate the installation of this script on the router. You must install the script manually.

Analytics Information on the Topology Map

You can view the following analytics information on the topology map:

- Delay A-Z
- Measured delay A-Z
- Interface utilization
- RSVP bandwidth A-Z
- RSVP utilization A-Z
- RSVP live utilization A-Z

To view the information on the topology map, right-click on a link and select **Link Label** > *Parameter* as shown in [Figure 86 on page 897](#).

Figure 86: Parameters to View Link Information

The screenshot displays the Paragon Automation Pathfinder Insights Planner interface. The main area shows a network topology map with nodes labeled vmx101 through vmx108 and iosx09. A context menu is open over a link between vmx105 and vmx108. The menu options are:


- Grouping >
- Select All Nodes
- Show All Nodes and Links
- Save to Default Map Layout
- Layout >
- Node Label >
- Link Label >
- Favorites > Link Label
- Subview >
- Timeline
- Pause Network Event Processing
- Reload Network

On the right side, a list of parameters is visible, including:

- Name
- Node Name A::Z
- Interface A::Z
- IP A::Z
- IP, SID A::Z
- TE Metric A::Z
- Bandwidth A::Z
- Delay A::Z
- Interface Util A::Z
- ISIS1 Metric A::Z
- ISIS2 Metric A::Z
- Measured Delay A::Z
- OSPF Metric A::Z
- Packet Loss A::Z
- RSVP Bandwidth A::Z
- RSVP Util A::Z
- RSVP Live Util A::Z
- SID A::Z
- Slices
- SRLG
- TE Admin Group A::Z
- Hide Label

[Figure 87 on page 898](#) shows the measured delay on a link displayed on the topology map.

Figure 87: Measured Delay Displayed on a Link



You can similarly view interface utilization, bandwidth, and other parameters of a link by right-clicking on a link and selecting the respective parameter for Link Label.

Monitor Workflows

IN THIS CHAPTER

- [Action Engine Workflow Overview | 899](#)
- [About the Workflows Monitor Page | 900](#)
- [About the Workflows Page | 902](#)
- [Manage Action Engine Workflows | 903](#)

Action Engine Workflow Overview

NOTE: Action Engine Workflow is a Beta feature.

When creating rules, Paragon Insights includes the ability to run user-defined actions (UDAs) as part of a trigger. UDAs are Python scripts that can be configured to be triggered by a Paragon Insights rule. For more information on UDAs, see "[Understand User-Defined Actions and Functions](#)" on page 593.

You cannot track and manage UDAs, pause an action, retry a failed action, and resume a paused action. However, Paragon Automation supports action engine workflow monitoring. An action engine workflow is an action engine that you can use to configure a set of tasks (instances). You can configure action engine workflows, monitor existing action engine workflows, and manage action engine workflow instances from the Paragon Automation GUI. See "[Manage Action Engine Workflows](#)" on page 903.

You can view action engine workflows that you created by using the CLI in the Paragon Automation GUI. You can perform the following actions from the CLI:

- run NETCONF command
- run Arbitrary executable files such as Python, Bash, Ruby
- run a command to send notification messages

RELATED DOCUMENTATION

[About the Workflows Monitor Page | 900](#)

[Understand User-Defined Actions and Functions | 593](#)

About the Workflows Monitor Page

IN THIS SECTION

- [Tasks You Can Perform | 900](#)
- [Field Descriptions | 901](#)

NOTE: Action Engine Workflow is a Beta feature.

To access the **Workflows Monitor** page from the Paragon Automation graphical user interface (GUI), click **Monitoring > Action Engine**. An action engine workflow is an action engine that you can use to configure a set of tasks (instances).

The **Workflows Monitor** page lists the various action engine workflows running in Paragon Automation.

Tasks You Can Perform

You can perform the following tasks from the **Workflows Monitor** page.

- Run an action engine workflow.
- Resume a suspended instance.
- Stop an instance.
- Filter instances.

For more information, see "[Manage Action Engine Workflows](#)" on page 903.

Field Descriptions

Field	Description
Workflow Name	Displays the name of the action engine workflow.
Total	Displays the total number of instances in the action engine workflow.
Running	Displays the total number of instances that are currently running the action engine workflow.
Suspended	Displays the total number of instances that are suspended.
Failed	Displays the total number of failed instances.
Completed	Displays the total number of completed instances in the action engine workflow.
Run Workflow	Click this button to run the action engine workflow(s) that you have selected.
Arguments	Displays the list of predefined arguments.
Additional Arguments	Displays the list of arguments that you have added.
Filter (icon)	Click this icon to set filters for instances within an action engine workflow.

RELATED DOCUMENTATION

[Action Engine Workflow Overview | 899](#)

[Manage Action Engine Workflows | 903](#)

About the Workflows Page

IN THIS SECTION

- [Tasks You Can Perform | 902](#)
- [Field Descriptions | 902](#)

NOTE: Action Engine Workflow is a Beta feature.

To access **Workflows** page from the Paragon Automation graphical user interface (GUI), click **Configuration > Action Engine**. An action engine workflow is an action engine that you can use to configure a set of tasks (instances). You can add new action engine workflows and delete existing action engine workflows from this page.

Tasks You Can Perform

- Configure an action workflow.
- Delete an action workflow.

For more information, see "[Manage Action Engine Workflows](#)" on page 903.

Field Descriptions

Field	Description
Workflow Name	Displays the name of the action engine workflow.
Description	Displays a brief description of the action engine workflow.
(+) icon	Click this icon to configure a new action engine workflow.

(Continued)

Field	Description
Filter icon	Click this icon to set filters for action engine workflows.

RELATED DOCUMENTATION

[Action Engine Workflow Overview | 899](#)
[Manage Action Engine Workflows | 903](#)
Manage Action Engine Workflows**IN THIS SECTION**

- [Add an Action Engine Workflow | 903](#)
- [Run an Action Engine Workflow | 907](#)
- [Stop an Instance | 908](#)
- [Resume a Suspended Instance | 908](#)
- [Filter Instances | 909](#)
- [Delete an Action Engine Workflow | 909](#)

NOTE: Action Engine Workflow is a Beta feature.

You can configure action engine workflows, and manage existing action engine workflows from the Paragon Automation GUI. You can add or edit action engine workflow commands, conditions, inputs, and outputs while creating an action engine workflow. You can also create action engine workflows by using the CLI. See "[Action Engine Workflow Overview](#)" on page 899.

Add an Action Engine Workflow

Follow these steps to add an action engine workflow:

1. Click **Configuration > Action Engine** .

The **Workflows** page appears.

2. Click the plus (+) icon to add an action engine workflow.

The **Add New Workflow** page appears. The **General** tabbed page appears by default.

3. Enter the following information in the **General** tabbed page:

- a. Enter a name for the action engine workflow in the **Name** text box.

- b. Enter a description for the action engine workflow in the **Description** text box.

- c. Select an entry task from the **Entry Task** drop-down list.

You must add a task before you can select the task from the **Entry Task** drop-down list. To add a task, see Step 4.

An entry task is the first task that is executed when you run an action engine workflow.

- d. Select an exit task from the **Exit Task** drop-down list.

You must add a task before you can select the task from the **Exit Task** drop-down list. To add a task, see Step 4.

An exit task is the last task (for example, a clean up task) that is executed at the end of an action engine workflow sequence.

4. Click **Tasks** to view the **Task** tabbed page.

5. Click (+) to add a task.

Enter the following information:

- a. Enter a name for the task in the **Name** text box.

- b. Enable or disable the **Parallel** toggle button.

Enable the **Parallel** toggle button to run all steps in a task simultaneously.

- c. Click the (+) icon to add a new step.

A row is added to the **Steps** section.

In the row that is added:

- a. Enter a name for the step in the **Name** text box.

- b. Enter a description for the step in the **Description** text box.

- c. Select dependencies from the **Dependencies** drop-down list.

- d. Select an action type from the **Action Type** drop-down list.

Follow these steps to add a command:

- a. Click **Edit Commands** to add a new command.

The **ADD/EDIT COMMANDS** pop-up appears.

- b. Click **+Add New Command** to add a new command.
The **New Command** section appears in the **ADD/EDIT COMMANDS** pop-up.
- c. Enter a value for the command tag in the **Command Tag** text box.
- d. Enter a value in the **Commands** list box. You can select more than one command.
Click **X** to remove the command that you selected.
- e. Enter a value in the **Arguments** list box. You can select more than one argument.
Click **X** to remove the argument that you selected.
- f. Enter a value in the **Device** list box. You can select more than one device.
Click **X** to remove the device that you selected.
- g. Enter a value in the **Device Group** list box. You can select more than one device group.
Click **X** to remove the device group that you selected.
- h. Enter a value in the **Environment** list box.
- i. Select an output from the **Output Type** list box.
- j. Enable or disable the **Ignore** toggle button.
You can enable the **Ignore** button to ignore steps.
- k. Set repeat parameters in the **Repeat** field.
You can determine if you want to repeat a failed step or not.
- l. The default delay value displayed is 10 seconds.
After you have set repeat parameters to repeat a step that has failed, there is a delay of 10 seconds before the step is repeated again.
- m. Click **OK** to confirm.
The new command is added.

Follow these steps to add a condition:

- a. Click **Edit Conditions** to add new conditions.
The **ADD/EDIT CONDITIONS** pop-up appears.
- b. Enter the conditions in the **Conditions** text box.
- c. Select a condition type from the **Conditions Type** list box.
- d. Enter a description for the condition in the **Condition Description** text box.
- e. Click **OK** to confirm.

The new condition is added.

Follow these steps to add inputs:

- a. Click **Edit Inputs** to add new inputs.
The **ADD/EDIT INPUTS** pop-up appears.
- b. Click the (+) icon to add new input.
- c. Enter a name for the input in the **Name** field.
- d. Enter a value for the input in the **Value** field.
- e. Click **OK** to confirm.
The operation is successful message is displayed in the **ADD/EDIT INPUTS** pop-up.
- f. Click **Close** to close the **ADD/EDIT INPUTS** pop-up.

Follow these steps to add outputs:

- a. Click **Edit Output** to add new outputs.
The **ADD/EDIT OUTPUT** pop-up appears.
- b. Click the (+) icon to add new input.
- c. Select a name for the output from the **Name** list.
- d. Enter a description for the output in the **Description** text box.
- e. Enter a value for the command tag in the **Command Tag** field.
- f. Select output type from the **Output Type** list box. See [Table 139 on page 906](#).
- g. The field displayed depends on the output type that you have selected. See [Table 139 on page 906](#).

Table 139: Output Type and Corresponding Fields

Output Type	Field
Grok	Pattern
XML	XPath
JSON	JQ Path
Artifact	Path

Table 139: Output Type and Corresponding Fields (Continued)

Output Type	Field
Regex	Pattern
Result	

- h. Click **OK**> to confirm.
The **operation is successful** message is displayed in the **ADD/EDIT OUTPUT** pop-up.
 - i. Click **Close** to close the **ADD/EDIT OUTPUT** pop-up.
 - j. Click the **✓** icon to add this row to the **Steps** section.
6. Click **Arguments** tab.
 7. On the Arguments tabbed page, click the plus (+) icon to add a new argument.
 - a. Enter a name for the argument in the **Name** text box.
 - b. Click **Ok** to confirm.
 8. Do any one of the following:
 - a. Click **Save** to save the action engine workflow.
 - b. Click **Save & Deploy** to save and deploy the action engine workflow.
You have now added and deployed an action engine workflow. To monitor the action engine workflows that you have added, see **Monitoring > Action Engine**.

Run an Action Engine Workflow

After you add an action engine workflow, you can run the action engine workflow by following these steps:

1. Click **Monitoring > Action Engine**.
The **Workflows Monitor** page appears.
2. Select the action engine workflow you want to run by selecting the check box next to the name of the action engine workflow.
3. Click **Run Workflow**.
The **Run Workflow <workflow name>** pop-up appears.
4. In the **Run Workflow <workflow name>** pop-up that appears, you can:
 - a. View the list of preconfigured arguments for the action engine workflow.
 - b. Configure additional arguments.
To configure additional arguments, click (+).

The **Additional Arguments** fields that you can configure are displayed.

- i. Enter a name in the **Name** text box to identify this additional argument.
The name you enter must be in the `[a-zA-Z][a-zA-Z0-9_-]*$` regular expression format. This format states that the first character of the name can start with a-z or A-Z. The name cannot start with a number or a special character. However, you can use numbers, `_`, and `-` within the name.

The maximum length is 64 characters.
- ii. Select an additional argument type from the **Type** drop-down list.
Available options: string, list, password, device, device-group, network-group
- iii. Select a value from the options available.
The options you can choose from depend on the additional argument **Type** that you have selected.

You can add one or more than one arguments.

5. Click **OK** to confirm settings and to run the action engine workflow.

Stop an Instance

You can stop an instance that is currently running.

To stop an instance:

NOTE: You cannot resume (restart) an instance that you have stopped.

1. Click **Monitoring > Action Engine**.
The **Workflows Monitor** page appears.
2. Click an action engine workflow to view the instances listed under it.
3. Select the instance that is currently running by selecting the check box next to the name of the instance.
4. Click **Stop** to stop the instance.

Resume a Suspended Instance

To resume a suspended instance:

1. Click **Monitoring > Action Engine**.
The **Workflows Monitor** page appears.
2. Click an action engine workflow to view the instances listed under it.
3. Select a suspended instance, by selecting the check box next to the name of the instance.
4. Click **Resume** to restart the instance.

Filter Instances

To filter instances within an action engine workflow:

1. Click **Monitoring > Action Engine**.
The **Workflows Monitor** page appears.
2. Click **Filter**, and then click **Add Filter** from the **Filter** drop-down list.
The **Add Criteria** pop-up appears.
3. Enter the following information in the **Add Criteria** pop-up.
 - a. Select the field that you want to apply the filter to, from the **Field** drop-down list.
 - b. Select the conditions that you want to apply to the field, from the **Condition** drop-down list.
 - c. Enter the start and finish time that you want to apply to the filter, in the **Value** text box.
4. Click **Add** to apply the filter.

Delete an Action Engine Workflow

To delete an action engine workflow:

1. Click **Configuration > Action Engine**.
The **Workflows** page appears.
2. Select the action engine workflow you want to delete by selecting the check box next to the name of the instance.
3. Click the **Delete** icon.
The **Delete Workflow** pop-up appears.
4. In the **Delete Workflow** pop-up that appears, click **Ok** to delete the action engine workflow.

10

PART

Reports

Health Reports | 911

Network Reports | 914

Maintenance Reports | 918

Inventory Reports | 934

Demand Reports | 947

Health Reports

IN THIS CHAPTER

- [About the Health Reports Page | 911](#)
- [View and Download Health Reports | 912](#)
- [Compare Differences in Health Reports | 913](#)

About the Health Reports Page

IN THIS SECTION

- [Tasks You Can Perform | 911](#)
- [Field Descriptions | 912](#)

To view this page, go to **Reports > Health Reports**.

You can view health reports generated for device groups and network groups on the **Health Reports** page. These reports include alarm statistics, device or network health data, as well as device-specific information (such as hardware and software specifications).

The reports displayed in this page are configured through report settings, scheduler settings, and destination settings on the **Ingest Settings** page.

You can click **Refresh** on the **Health Reports** page to refresh the page and to display the latest health reports. You can also enable **Auto Refresh** and specify the refresh rate in seconds to automatically load new reports.

Tasks You Can Perform

You can perform the following tasks on this page:

- View and download a report. See "[View and Download Health Reports](#)" on page 912.
- Generate a diff report. See "[Compare Differences in Health Reports](#)" on page 913

Field Descriptions

[Table 140 on page 912](#) describes the fields on the Health Reports.

Table 140: Fields on the Health Reports Page

Fields	Description
Report Id	Unique ID generated automatically for a report.
Group Name	Name of the device and network groups to which the report is applied.
Report Name	Name of the report, which you configure in the report profile.
Scheduler Name	Name of the scheduler(s) selected in the report profile.
Generated On	The day, date, month, and time (24-hour format) when the report is generated.

RELATED DOCUMENTATION

| [Configure Report Settings](#) | 628

View and Download Health Reports

You can view the latest health reports on the **Reports > Health Reports** page. A health report provides information on device and network health, device-related hardware and software specifications, and alarm statistics,

You can view reports in the Health Reports page only if the destination type for a report is set to **disk**. If the destination type for the report is set to **email**, check the email inbox of the specified account for the report and open the attachment.

Ensure that you have configured the necessary ["destination settings" on page 633](#), ["scheduler settings" on page 630](#), and ["general report settings" on page 628](#) before you view and download health reports.

The reports in the Health Reports page are organized by the date and time at which they were generated. The most recent report is listed at the top of the table.

To find and download a report:

1. Search for a report within a column using the text box under the column heading.

You can search for a report by report Id, group name (device or network group name), report name, scheduler name, or the date a report was generated on.

While you search for a report, you can also click any column heading to sort reports based on the column category. To adjust the number of reports displayed, click the number of rows displayed at the bottom of the page and select a number.

2. Click name of the report to download it to your system.

RELATED DOCUMENTATION

[About the Health Reports Page | 911](#)

Compare Differences in Health Reports

When you generate two or more health reports, you can compare any two reports in a 'diff' report. A diff health report shows changes that you can use to determine device or network health, and suggest ways to troubleshoot if necessary.

To generate a diff health report:

1. Select **Reports > Health Reports**.
2. Select two reports that you want to compare, and click **Diff Report**.

The diff health report opens in a new browser tab.

Ensure that pop-ups are not blocked in your browser.

RELATED DOCUMENTATION

[View and Download Health Reports | 912](#)

[About the Health Reports Page | 911](#)

Network Reports

IN THIS CHAPTER

- [Network Reports Overview | 914](#)
- [View Network Reports | 915](#)

Network Reports Overview

Network reports help you identify network issues as a result of configuration errors and LSP discrepancies.

Network reports are of two types:

- Integrity check reports—To generate these reports, you must first run a device collection task with **Configuration** as a collection option to collect the router's configuration.
- LSP discrepancy reports—These reports are generated automatically when the topology is reloaded or when a PCEP session on a node flaps (in which case, the report is generated only for that node).

[Table 141 on page 914](#) explains the types of network reports. For details about each type of network report, see "[View Network Reports](#)" on page 915.

Table 141: Types of Network Reports

Report Type	Description
Integrity Check	Lists potential configuration errors in the router's configuration file as a result of the integrity checks performed by the Path Computation Server (PCS).
LSP Discrepancy	Lists details of the PCC-initiated and PCC-delegated LSPs that might require reprovisioning due to LSP discrepancies discovered by the Path Computation Server (PCS) when the topology is reloaded or when a PCEP session on a node flaps.

RELATED DOCUMENTATION

| [About the Health Reports Page](#) | 911

View Network Reports

Paragon Automation enables you to generate and view Network reports that help you identify network issues. For information on the types of Network reports and how to generate them, see "[Network Reports Overview](#)" on page 914.

To access the page for a specific network report, select **Reports > Network > Report-Name**.

You can perform the following actions on this page:

- View the following Network reports:
 - Integrity Check reports, see [Table 142 on page 916](#).
 - LSP Discrepancy reports, see [Table 143 on page 917](#).
- Download a report—Hover over the Download button on the top-right corner of the report and select either **CSV** or **JSON** to download the report as a comma-separated values file (CSV) or JavaScript Object Notation (JSON) file, respectively.
- View router configuration files—Select an item in the Integrity Check report and hover over the More button on the top-right corner of the report to select **Show Config**. Alternatively, hover over an item and click the details icon that appears.

The Configuration pane appears on the right side of the page, displaying the configuration for the selected item.

- Show or hide columns in a report—Click the vertical ellipsis icon at the top right corner of the report and select the columns that you want displayed in the report.

Only the columns that you selected are displayed.

[Table 142 on page 916](#) explains the fields displayed in the Integrity Check reports.

Table 142: Integrity Check Reports

Field	Description
Category	Category to which the integrity check belongs. Example: IS-IS, Tunnel, RSVP
Message	Description of the error. This column may also contain references to specific configuration files or host names when more than one configuration file is involved.
Detail	Details of the entity in which the error was detected.
Severity	Severity level (High, Medium, Low, and Warning) of the error: <ul style="list-style-type: none"> • High—Indicates a severe problem that may cause major issues in the network. • Medium—Indicates that the problem is not severe but should be fixed to prevent issues in the network. • Low—Indicates that the problem is not severe but should be fixed to prevent issues in the network. An error with Medium severity level has higher priority over an error with Low severity level. • Warning—Indicates a potential problem in the network and must be examined.
Error Source	Device in which the error was detected.
Source File	File in which the error was detected. Displays a configuration file name, or a description when more than one configuration file is involved.
Line Number	Line in the configuration file where the error was detected.
Line Content	Content of the line in the configuration file where the error was detected.

Table 142: Integrity Check Reports (Continued)

Field	Description
Msg ID	Index value assigned to the error for identification. NOTE: This field is reserved for future use.

[Table 143 on page 917](#) explains the fields displayed in the LSP Discrepancy reports.

Table 143: LSP Discrepancy Reports

Field	Description
Out-of-Sync Time	Date (in MM/DD/YYYY format) and time (in HH:MM:SS 24-hour or AM/PM format) at which the LSP discrepancy report was generated.
LSP Name	Name of the LSP that is impacted.
Node	Source Node for the LSP that is impacted.
Control Type	Control type of the LSP (PCE-initiated, Delegated, or Device-controlled) that is impacted.
Path Type	Path type of the LSP (Primary, Secondary, or Standby) that is impacted.
Description	Details of the discrepancy in the LSP. Example: Host is not pcep enabled

RELATED DOCUMENTATION

[About the Health Reports Page](#) | 911

Maintenance Reports

IN THIS CHAPTER

- [Maintenance Reports Overview | 918](#)
- [View Maintenance Reports | 919](#)

Maintenance Reports Overview

When an exhaustive failure maintenance simulation completes, several maintenance reports are generated which can be viewed from **Reports > Maintenance > Report-Name** page. For more information on each maintenance report, see "[View Maintenance Reports](#)" on page 919.

NOTE: If you have multiple maintenance event simulations active at a time, you can choose to view reports for a particular event by selecting the event name from the **Branches** list at top-left corner of the page.

[Table 144 on page 918](#) describes the reports generated after an exhaustive maintenance simulation.

Table 144: Types of Maintenance Reports

Report Name	Description
Link Oversubscription	Lists the links that reached over 100% utilization during the exhaustive failure simulation.
Link Utilization Changes	Shows changes in link RSVP bandwidth reservation if all LSPs were to be routed over their optimized paths instead of their current paths.

Table 144: Types of Maintenance Reports (Continued)

Report Name	Description
LSP Path Changes	Shows changes to the tunnel paths, number of hops, path cost, and delay for PCE-initiated and PCC-delegated LSPs as a result of path analysis.
Maintenance Simulation	Shows link utilization and LSP routing changes caused by maintenance events during failure simulation.
Path Delay	Shows the worst path delay and distance experience by each tunnel and the associated failure event that caused the worst-case scenario.
Peak Interface Utilization	Shows physical interfaces report with normal utilization, the worst utilization, and the causing events during exhaustive failure simulation.
Peak Link Utilization	Shows the peak utilization encountered from one or more elements that failed for each link.
Peak Simulation Summary	Shows the summary view of the count, bandwidth, and hops of the tunnels that were impacted or failed during simulation.
Peak Tunnel Failure	Lists the tunnels that were unable to reroute and the events that prevented the tunnels reroute during exhaustive failure simulation.

RELATED DOCUMENTATION

| [Simulate a Maintenance Event | 781](#)

View Maintenance Reports**IN THIS SECTION**

● [Tasks You Can Perform | 920](#)

- [Link Oversubscription Report | 921](#)
- [Link Utilization Changes Report | 922](#)
- [LSP Path Changes Report | 923](#)
- [Maintenance Simulation Report | 924](#)
- [Path Delay Report | 925](#)
- [Peak Interface Utilization Report | 926](#)
- [Peak Link Utilization Report | 928](#)
- [Peak Simulation Summary Report | 930](#)
- [Peak Tunnel Failure | 932](#)

To access an individual maintenance report, click **Reports** > **Maintenance** > *Report-Name*.

NOTE: If you have multiple maintenance event simulations active at a time, you can choose to view reports for a particular event by selecting the event name from the **Branches** list at top-left corner of the page.

Tasks You Can Perform

You can perform the following tasks on individual report page:

- View details about individual reports.
 - Link Oversubscription, see [Table 145 on page 921](#).
 - Link Utilization Changes, see [Table 146 on page 922](#).
 - LSP Path Changes, see [Table 147 on page 923](#).
 - Maintenance Simulation, see [Table 148 on page 924](#).
 - Path Delay, see [Table 149 on page 925](#).
 - Peak Interface Utilization, see [Table 150 on page 926](#).
 - Peak Link Utilization, see [Table 151 on page 928](#).
 - Peak Simulation Summary, see [Table 152 on page 931](#).

- Peak Tunnel Failure, see [Table 153 on page 932](#).
- Download the report—Hover over the **Download** button and select a format, CSV or JSON, to export the detailed maintenance report.
- Show/Hide Columns—Choose to show or hide a specific column in the table.

Hover over the **More Options** (vertical ellipsis) > **Show/Hide Columns** and select the *Column-Name* check box of the columns you want to display on the page.

Link Oversubscription Report

If link oversubscription does not occur for the maintenance events, the report will be empty.

[Table 145 on page 921](#) describes the fields on the Link Oversubscription Page.

Table 145: Fields on the Link Oversubscription Page

Field	Description
Simulation Type	Type of exhaustive maintenance simulation. It can be Node, Link, SRLG, or a combination of the three types.
Event	Name of the network element under maintenance. It can be node name, link name or facility (SRLG) name.
Layer	Type of Layer—Demand or Tunnel.
Link Name	Name of the link.
Node A	Name or IP address of Node A (ingress node).
Interface	Name of interface A.
Node Z	Name or IP address of Node Z (egress node).
Change	Change in the link bandwidth (in bps).
Bandwidth	Bandwidth (in bps) of the link.

Table 145: Fields on the Link Oversubscription Page (Continued)

Field	Description
Used Bandwidth	Bandwidth (in bps) allocated for tunnels resulting from the specified maintenance event simulation.

Link Utilization Changes Report

[Table 146 on page 922](#) describes the fields on the Link Utilization Changes Page.

Table 146: Fields on the Link Utilization Changes Page

Field	Description
Link Name	Name of the link.
Node:Interface	Node and interface name of the local node.
Remote Node	Name of the node at the remote end.
Change	Change made to the link. Link Down is displayed when the link is under maintenance. If there is a change in link utilization, this field will be empty.
Bandwidth	Total bandwidth (in bps) available for traffic.
New Utilization	Percentage value of new link utilization. Range is between 0 and 1. For example, 0.8 implies 80% link utilization.
Utilization	Percentage value of original link utilization. Range is between 0 and 1. If the link is under maintenance, this field will be empty.
Difference	Difference between the new and original link utilization.

LSP Path Changes Report

Table 147 on page 923 describes the fields on the LSP Path Changes page.

Table 147: Fields on the LSP Path Changes Page

Field	Description
LSP Name	Name of the LSP.
Node A	Name of Node A (ingress node).
Node Z	Name of Node Z (egress node).
Hop Count	Original hop count of the LSP.
New Hop Count	New hop count of the LSP resulting from the specified maintenance event simulation.
Path Cost	An integer value for the cost associated with the original LSP path. When an LSP is routed over multiple links, path cost is the cumulative value of all individual link costs that the LSP goes through.
New Path Cost	An integer value for the cost associated with the new (rerouted) LSP path.
Path	Displays the IP addresses that the original LSP goes through from Node A.
New Path	Displays the IP addresses that the new (rerouted) LSP goes through from Node A.
Protection	Displays the name of the standby or secondary path associated with the LSP, if any.
Delay	Delay (in milliseconds) associated with the original LSP. When an LSP is routed over multiple links, delay is the cumulative value of all individual link delays that the LSP goes through.
New Delay	Delay (in milliseconds) associated with the new (rerouted) LSP.

Table 147: Fields on the LSP Path Changes Page (Continued)

Field	Description
Delay Change	Percentage increase or decrease in the delay. It is calculated by: $(\text{Delay}-\text{New Delay}/\text{Delay}) * 100$.
Bandwidth	Bandwidth (in bps) of the LSP.
New Bandwidth	Bandwidth (in bps) after the LSP is rerouted.
Type	Type of LSP (delegated or device controlled).

Maintenance Simulation Report

[Table 148 on page 924](#) describes the fields on Maintenance Simulation Page.

Table 148: Fields on the Maintenance Simulation Page

Field	Description
Event ID	Name of the maintenance event.
Status	Status of the maintenance event execution— Completed-Pass or Completed-Fail (execution is completed but with errors).
Simulation Time	Time when you click on Simulate on the maintenance tab of the network information table.
Fail Count	Total number of path routing failures that occurred during the maintenance event simulation.
Oversubscription Count	Total number of links where oversubscription (more than 100%) occurred during maintenance event simulation.

Table 148: Fields on the Maintenance Simulation Page (Continued)

Field	Description
Simulation Type	Type of exhaustive maintenance simulation. It can be Node, Link, SRLG, or a combination of the three types.

Path Delay Report

[Table 149 on page 925](#) describes the fields on the Path Delay page.

Table 149: Fields on the Path Delay Page

Field	Description
Path Name	Name of the path.
From	Name of the ingress node.
To	Name of the egress node.
Bandwidth	Bandwidth (in bps) associated with the path.
Priority	Priority for the path traffic. Range 1 through 7.
Path	Displays the IP addresses that the LSP goes through.
Distance	Total path distance. NOTE: 0 indicates that path was not rerouted during the simulation.
Delay	Propagation delay of the path in normal mode.
Fail Count	Number of times this path was disconnected during the maintenance event failure simulation.

Table 149: Fields on the Path Delay Page (Continued)

Field	Description
Worst Distance	Worst path distance of the alternate routes that occurred during the maintenance event failure simulation.
Worst Delay	Worst propagation delay of the alternate routes that occurred during the maintenance event failure simulation.
Worst Delay Cause	<p>Cause of the worst delay during the maintenance event simulation. It can be due to one of the following:</p> <ul style="list-style-type: none"> • NDFAIL (node failure) • LINKFAIL (link failure) • FACFAIL (facility failure) • SIMPLACE (alias for element that is put under maintenance)
Delay Cause Event	Network element that caused the worst delay during the maintenance event simulation. It can be caused by a node, a link (From and To nodes are displayed), or a facility.

Peak Interface Utilization Report

[Table 150 on page 926](#) describes fields on the Peak Interface Utilization page.

Table 150: Fields on the Peak Interface Utilization Page

Field	Description
Node:Interface	Name of the node and interface associated with the interface.
Link Count	Number of links associated with the interface.
Bandwidth	Total bandwidth (in bps) available for the user traffic.

Table 150: Fields on the Peak Interface Utilization Page *(Continued)*

Field	Description
Used Bandwidth	Bandwidth (in bps) used by tunnels in normal mode.
Peak Bandwidth	Maximum bandwidth (in percentage) used by tunnels during the maintenance event failure simulation. For example, 5 implies 50% utilization.
Utilization	Bandwidth utilization (in percentage). It is calculated as $100 * (\text{UsedBandwidth} / \text{TotalBandwidth})$.
Peak Utilization	Peak bandwidth utilization (in percentage). It is calculated as $100 * (\text{PeakBandwidth} / \text{TotalBandwidth})$.
Tunnel Count	Number of tunnels carried by the link in normal mode.
Peak Tunnel Count	Maximum number of tunnels carried by the link during the maintenance event failure simulation.
Worst Load Cause	<p>Cause of the worst load (bandwidth utilization) during the maintenance event simulation. It can be due to one of the following:</p> <ul style="list-style-type: none"> • NDFAIL (node failure) • LINKFAIL (link failure) • FACFAIL (facility failure) • SIMPLACE (alias for element that is put under maintenance)
Load Cause Event	Network element that caused the worst load during the maintenance event simulation. It can be caused by a node, a link (From and To nodes are displayed), or a facility.

Table 150: Fields on the Peak Interface Utilization Page (Continued)

Field	Description
Worst Tunnel Cause	<p>Cause of the worst (or increased) tunnel count during the maintenance event simulation. It can be due to one of the following:</p> <ul style="list-style-type: none"> • NDFAIL (node failure) • LINKFAIL (link failure) • FACFAIL (facility failure) • SIMPLACE (alias for element that is put under maintenance)
Tunnel Cause Event	<p>Network element that caused the worst tunnel count during the maintenance event simulation. It can be caused by a node, a link (From and To nodes are displayed), or a facility.</p>

Peak Link Utilization Report

[Table 151 on page 928](#) describes the fields on Peak Link Utilization page.

Table 151: Fields on the Peak Link Utilization Page

Field	Description
Link Name	Name of the link.
Node A:Interface	ID of node and associated interface for node A where the link originates.
Loc A	Name of the source node where the link originates (ingress).
Node Z:Interface	ID of node and associated interface for node Z where the link terminates (egress).
Loc Z	Name of the node at Z end (egress).

Table 151: Fields on the Peak Link Utilization Page (Continued)

Field	Description
Vdr	The vendor associated with this link. Possible values for vendors include those that are specific to a certain country or region, and are listed in the tariff database. If a vendor is not specified, this value is set to the default value DEF.
Link Type	The type of link being used. The trunk type is subsequently used in determining link pricing and bandwidth availability.
Bandwidth	Total bandwidth (in bps) available for the user traffic (between node A and Z).
Used Bandwidth	Bandwidth (in bps) used by tunnels in normal mode on this interface.
Peak Bandwidth	Maximum bandwidth (in bps) used by tunnels during the maintenance event failure simulation (between node A and Z).
Utilization	Bandwidth utilization (in percentage). It is calculated as per the following formula: $100 * (\text{UsedBandwidth} / \text{TotalBandwidth})$
Peak Utilization	Peak bandwidth utilization (in percentage). It is calculated as per the following formula: $100 * (\text{PeakBandwidth} / \text{TotalBandwidth})$
Tunnel Count	Number of tunnels carried by the link in normal mode.
Peak Tunnel Count	Maximum number of tunnels carried by the link during the maintenance event failure simulation.
Oversubscription Count	Number of failures that caused used bandwidth to exceed $(1 - \text{fatpct}) * \text{TotalBw}$ where $\text{fatpct} = 0.00\%$.

Table 151: Fields on the Peak Link Utilization Page *(Continued)*

Field	Description
Worst Load Cause	<p>Cause of the worst load (bandwidth utilization) during the maintenance event simulation. It can be due to one of the following:</p> <ul style="list-style-type: none"> • NDFAIL (node failure) • LINKFAIL (link failure) • FACFAIL (facility failure) • SIMPLACE (alias for element that is put under maintenance)
Load Cause Event	<p>Network element that caused the worst load during the maintenance event simulation. It can be caused by a node, a link (From and To nodes are displayed), or a facility.</p>
Worst Tunnel Cause	<p>Cause of the worst (or increased) tunnel count during the maintenance event simulation. It can be due to one of the following:</p> <ul style="list-style-type: none"> • NDFAIL (node failure) • LINKFAIL (link failure) • FACFAIL (facility failure) • SIMPLACE (alias for element that is put under maintenance)
Tunnel Cause Event	<p>Network element that caused the worst tunnel count during the maintenance event simulation. It can be caused by a node, a link (From and To nodes are displayed), or a facility.</p>

Peak Simulation Summary Report

[Table 152 on page 931](#) describes the fields on Peak Simulation Summary page.

Table 152: Fields on the Peak Simulation Summary Page

Field	Description
Simulation Type	Type of exhaustive maintenance simulation. It can be Node, Link, SRLG, or a combination of the three types.
Event	Name of the network element under maintenance. It can be node name, link name or facility (SRLG) name.
UP/Down	Operation performed in the simulation—Up or Down.
Layer	Type of Layer—Demand or Tunnel.
Impact Count	Number of tunnels impacted by the simulation.
Impact Bandwidth	Total bandwidth (in bps) of the impacted demand or tunnel.
Fail Count	Number of disconnected flows (tunnels that are terminated at failed nodes are not included).
Fail Bandwidth	Total bandwidth (in bps) of disconnected flows
Fail Bandwidth Percentage	$100 * \text{FailedBandwidth} / \text{TotalFlowBandwidth}$ percentage
Highest Priority Fail	Highest priority of failed flows.
Oversubscription Count	Number of links where bandwidth oversubscription has occurred.
Max Hop	Maximum path hop count after failure.
Average Hop	Average path hop count after failure.
Terminated Count	Number of flows terminated at failed nodes.

Table 152: Fields on the Peak Simulation Summary Page (Continued)

Field	Description
Terminated Bandwidth	Total bandwidth (in bps) of flows terminated at failed nodes.

Peak Tunnel Failure

[Table 153 on page 932](#) describes the fields on the Peak Tunnel Failure page.

Table 153: Fields on the Peak Tunnel Failure Page

Field	Description
Simulation Type	Type of exhaustive maintenance simulation. It can be Node, Link, SRLG, or a combination of the three types.
Event	Name of the network element under maintenance. It can be node name, link name or facility (SRLG) name.
Layer	Type of Layer—Demand or Tunnel.
Path Name	Name of the path.
From	Name of the ingress node.
To	Name of the egress node.
ToIPAddr	IP address of the egress (To) node.
Bandwidth	Total bandwidth (in bps) available for the user traffic.
Priority	Setup and Hold priority for the LSP traffic. Range is 0 (highest priority) through 7 (lowest priority). The default is 7, which is the standard MPLS LSP definition in Junos OS. For example, 7,0 implies that 7 is the setup priority and 0 is the hold priority.

Table 153: Fields on the Peak Tunnel Failure Page (*Continued*)

Field	Description
Path Comment	<p>Comment (in text format) about the failed tunnel path. It can have one of the following values:</p> <ul style="list-style-type: none"> • Not Routed—Tunnel has failed and has not been routed. • Time Expired—Tunnel is no longer active in the network. For example, when you schedule an LSP, the LSP will be active only during the scheduled time period. When you run a maintenance simulation when this LSP is not in the scheduled time period, it will be marked here as Time Expired.
Info	<p>Only when a node failure occurs during maintenance simulation, this field displays one of the following values:</p> <ul style="list-style-type: none"> • Pass Through—Implies that this LSP passes through but does not originate or terminates at the failed node. • Terminated at Node—Implies that this LSP originates or terminates on the failed node.

Inventory Reports

IN THIS CHAPTER

- [Inventory Reports Overview | 934](#)
- [View Inventory Reports | 936](#)

Inventory Reports Overview

Inventory reports provide details of the hardware inventory (such as devices, device parts, line cards, and so on) in the network.

[Table 154 on page 934](#) explains the different types of inventory reports. For details about each type of inventory report, see ["View Inventory Reports" on page 936](#).

To generate inventory reports, you must first:

1. Run a device collection task with **Equipment CLI data** as the collection option to collect equipment CLI data. See ["Add a Device Collection Task" on page 978](#).
2. Run a network archive task with **Process Equipment CLI** selected to parse the equipment CLI data collected during the device collection task. See ["Add a Network Archive Task" on page 990](#).

Table 154: Types of Inventory Reports

Report Type	Description
Devices	Lists the details of devices present in the network. The details include hostname, chassis model, OS version installed, and so on.

Table 154: Types of Inventory Reports (Continued)

Report Type	Description
Device Parts	<p>Lists the details of parts present in the devices in the network.</p> <p>The parts include the chassis, Flexible PIC Concentrators (FPCs), fan trays, and power modules present in the different devices in the network. The details include the part name, serial number, device to which the part belongs, vendor of the part, cost of the part, and so on.</p>
Line Cards	<p>Lists the details of line cards present in the devices in the network.</p> <p>The details include the card name, name of the device in which the line card is present, IP address, connected ports, reserved ports, shutdown ports, vendor, and so on.</p>
Line Cards Usage	<p>Lists the details pertaining to the usage of line cards present in the devices in the network.</p> <p>The details include the line card name, name of the device in which the line card is present, card ID, and so on.</p>
Miscellaneous Parts	<p>Lists the miscellaneous parts of devices in the network.</p> <p>The miscellaneous parts include midplane, power entry modules, control plane, and so on. The details include the name of the part, name of the device in which the part is present, card ID, part vendor, IP address, and so on.</p>
Physical Interfaces	<p>Lists the details of physical interfaces present in the devices in the network.</p> <p>The details include the interface name, name of the device in which the physical interface is present, interface vendor, card ID, and so on.</p>
Transceivers	<p>Lists the details of transceivers present in the devices in the network.</p> <p>The details include the transceiver name, name of the device in which the transceiver is present, vendor, IP address, serial number, and so on.</p>

RELATED DOCUMENTATION

| [About the Health Reports Page](#) | 911

View Inventory Reports

Paragon Automation enables you to generate and view Inventory reports that provide details of the hardware inventory in the network. For information on the types of Inventory reports and how to generate them, see ["Inventory Reports Overview" on page 934](#).

To access the page for to a specific Inventory report, select **Reports > Inventory > Report-Name**.

You can perform the following actions on this page:

- View the following Inventory reports:
 - Devices, see [Table 155 on page 936](#).
 - Device Parts, see [Table 156 on page 939](#).
 - Line Cards, see [Table 157 on page 940](#).
 - Line Cards Usage, see [Table 158 on page 942](#).
 - Miscellaneous Parts, see [Table 159 on page 943](#).
 - Physical Interfaces, see [Table 160 on page 944](#).
 - Transceivers, see [Table 161 on page 946](#).
- Download a report—Hover over the Download button on the top-right corner of the report and select either **CSV** or **JSON** to download the report as a comma-separated values file (CSV) or JavaScript Object Notation (JSON) file, respectively.
- Show or hide columns in a report—Click the vertical ellipsis icon at the top-right corner of the report and select the columns that you want displayed in the report.

Only the columns that you selected are displayed.

[Table 155 on page 936](#) explains the fields displayed in the Devices report.

Table 155: Devices Report

Field	Description
Device Name	Name of the device present in the network.
Vendor	Name of the device vendor. Example: Juniper

Table 155: Devices Report (Continued)

Field	Description
IP Address	IPv4 address of the device.
Source	Filename of the configuration file from where the device details are collected.
System Name	System name for the device.
Description	Description for the device.
Contact	Contact name for the device.
Location	Physical location of the device.
Last update by CLI	Date (in MM:DD:YYYY format) and time (in HH:MM:SS 12-hour or AM/PM format) at which the configuration file was last updated as a result of device collection.
Chassis Type	Type of chassis. Example: WS-C6509-E
Hardware Version	Hardware revision model of the chassis.
Hardware ID	Serial number of the chassis.
ROM Version	ROM monitor version.
ROM System Version	ROM system software version.
Configuration Register	SNMP configuration register value.
Memory	RAM (in bytes) available to the CPU.

Table 155: Devices Report (Continued)

Field	Description
Configuration Memory	Non-volatile configuration memory available, in bytes.
Configuration Memory In Use	Non-volatile configuration memory used, in bytes.
Boot Image	Information related to the boot image. Example: a string starting with flash:, bootflash:, disk0:
Processor	Processor used in the device. Example: PowerPC405
OS Version	Version of OS that is currently installed on the device.
Management IP Address	Management IP Address that is used by Paragon Planner to discover the device.
Model	Additional details not known from the vendor, for example, by distinguishing JUNIPER and JUNIPER_EX.
OS Family	OS family to which the device belongs. Example: JUNOS
Last Update by SNMP	Date (in MM:DD:YYYY format) and time (in HH:MM:SS 12-hour or AM/PM format) at which the configuration file was last updated as a result of SNMP collection.
Hostname	Hostname of the device.
IPv6	IPv6 address of the device, if any.
AS	Autonomous system (AS) number assigned to the device.

[Table 156 on page 939](#) explains the fields displayed in the Device Parts report.

Table 156: Device Parts Report

Field	Description
Name	Name of the device part.
Device Name	Name of the device to which this part belongs.
Vendor	Name of the device vendor. Example: Juniper
IP Address	IP address of the device to which the part belongs.
Part	Part number associated with the device part.
Serial Number	Serial number of the device part.
Description	Description for the device part.
Estimated Cost	Estimated cost of the device part.
Hostname	Hostname of the device.
AS	AS number assigned to the device.
Model	Additional details not known from the vendor, for example, by distinguishing JUNIPER and JUNIPER_EX.
Card ID	Location of the card on the device. Example: S-0/1

[Table 157 on page 940](#) explains the fields displayed in the Line Cards report.

Table 157: Line Cards Report

Field	Description
Card Name	Name of the line card. For example: WS-X6704-10GE
Device Name	Name of the device to which this part belongs.
Vendor	Name of the device vendor. Example: Juniper
IP Address	The IP address of the device to which the part belongs.
Card ID	Location of the card on the device. Example: S-0/1
Connected Ports	Sum of the number of ports that meet one or more of the following conditions: <ul style="list-style-type: none"> • Number of ports that have a corresponding link in the model. • Number of ports that are operationally up. • Number of ports that have an IP address defined on the interface.
Reserved Ports	Number of reserved ports in the line card. A port is considered reserved if the interface description in the configuration file contains one or more of the following keywords: <ul style="list-style-type: none"> • Reserved • Booked • Spare
Shutdown Ports	Number of ports that have operational status or administrative status as Down.

Table 157: Line Cards Report (Continued)

Field	Description
Port Count	Number of ports on the line card.
Part	Part number of the line card.
Serial number	Serial number of the line card.
FRU Line Cards	Name of the field-replaceable unit (FRU) line card (applicable only to Cisco devices).
FRU Route Memory	Route memory allocated to the FRU (applicable only to Cisco devices).
FRU Packet Memory	Packet memory allocated to the FRU (applicable only to Cisco devices).
L3 Engine Type	Type of L3 Engine (applicable only to Cisco devices).
L3 Engine	Details related to the L3 engine (applicable only to Cisco devices).
TAN	Top Assembly Number (TAN) or ordering number (applicable only to Cisco devices).
Description	Description for the line card.
Controller Memory	Controller memory (applicable only to Cisco devices).
Estimated Cost	Estimated cost of the line card.
Management IP Address	Management IP Address that is used by Paragon Planner to discover the device.
Version	Line card version information. Example: rev A0 ver 4
Hostname	Hostname of the device.

Table 157: Line Cards Report (Continued)

Field	Description
Operational State	Operational state of the device.
Admin State	Administrative state of the device.
AS	AS number assigned to the device.

[Table 158 on page 942](#) explains the fields displayed in the Line Cards Usage report.

Table 158: Line Cards usage Report

Field	Description
Card Name	Name of the line card. For example: WS-X6704-10GE
Device Name	Name of the device to which this part belongs.
Card ID	Location of the card on the device. Example: S-0/1
Connected Ports	Sum of the number of ports that meet one or more of the following conditions: <ul style="list-style-type: none"> • Number of ports that have a corresponding link in the model. • Number of ports that are operationally up. • Number of ports that have an IP address defined on the interface.
Reserved Ports	Number of reserved ports, based on the customizable keywords found in the description in the configuration file.
Shutdown Ports	Number of ports that have operational status or administrative status as Down.

Table 158: Line Cards usage Report (Continued)

Field	Description
Number of Ports	Number of ports on the line card.
Hostname	Hostname of the device.

[Table 159 on page 943](#) explains the fields displayed in the Miscellaneous Parts report.

Table 159: Miscellaneous Parts Report

Field	Description
Name	Name of the device part. Example: Power supply, fan, display
Device Name	Name of the device to which this part belongs.
Vendor	Name of the device vendor. Example: Juniper
IP Address	IP address of the device to which the part belongs.
Part	Part number associated with the device part.
Serial Number	Serial number of the device part.
Description	Description for the device part.
Estimated Cost	Estimated cost of the device part.
Hostname	Hostname of the device.
AS	AS number assigned to the device.

Table 159: Miscellaneous Parts Report (Continued)

Field	Description
Model	Additional details not known from the vendor, for example, by distinguishing JUNIPER and JUNIPER_EX.
Card ID	Location of the card on the device. Example: S-0/1

[Table 160 on page 944](#) explains the fields displayed in the Physical Interfaces report.

Table 160: Physical Interfaces Report

Field	Description
Interface	Name of the physical interface. Example: GigabitEthernet0/0/0/0
Device Name	Name of the device that contains the physical interface.
Vendor	Name of the device vendor. Example: Juniper
IP Address	IP address of the physical interface.
Admin Status	Administrative status (active, down, or unknown) of the physical interface.
IP Addr	Loopback IP address of the device, typically the first IP address in the IP address range is listed.
Media Type	Media type for the physical interface. Example: Ethernet, FE, GE
MTU	Maximum Transmission Unit (MTU) of the physical interface.

Table 160: Physical Interfaces Report (Continued)

Field	Description
Bandwidth	Bandwidth of the physical interface. Example: 1000.0M
Physical Address	Physical address (MAC address) for the physical interface.
Operational Status	Operational status (Active, Down, or Unknown) of the physical interface.
Description	Description for the physical interface.
IPv6	IPv6 address of the physical interface, if any.
Management IP Address	Management IP Address that is used by Paragon Planner to discover the device.
Card ID	Location of the card on the device. Example: S-0/1
Aggregated Link	Aggregated link (if any) associated with the physical interface. Example: PortChannel, Bundle-Ether
VLAN	List of VLANs on the physical interface. Example: 22 100-102 105-108
Switchport Mode	Switchport mode for layer 2 interfaces. Example: Access or Trunk
Hostname	Hostname of the device.
AS	AS number assigned to the device.

[Table 161 on page 946](#) explains the fields displayed in the Transceivers report.

Table 161: Transceivers Report

Field	Description
Name	Name of the transceiver. Example: SFP-10GBase-ER
Device Name	Name of the device.
Vendor	Name of the device vendor. Example: Juniper
IP Address	Loopback IP address of the device, typically the first IP address in the IP address range is listed.
Index	Location of the card on the device. Example: S-0/1
Part	The part number associated with the transceiver.
Serial Number	Serial number of the transceiver.
Estimated Cost	Estimated cost of the transceiver.
Hostname	Hostname of the device.
AS	AS number assigned to the device.
ContainedIn	Name of the line card that contains the transceiver.

RELATED DOCUMENTATION

| [About the Health Reports Page](#) | 911

Demand Reports

IN THIS CHAPTER

- [Demand Reports Overview | 947](#)
- [View Demand Reports | 948](#)

Demand Reports Overview

Demand reports provide information about network traffic. You can use this information to evaluate ingress and egress traffic patterns.

Demand reports are of four types:

- VPN Demand reports—To generate these reports, you must run the Demand Reports task, with VPN Demands selected as the report type.
- Group Demand reports—To generate these reports, you must first group nodes as a single entity in the topology map; see ["Group Nodes" on page 682](#). Then, you must run the Demand Reports task, with Group Demands selected as the report type.
- Label-switched paths (LSP) Demand reports—To generate these reports, you must run the Demand Reports task, with LSP Demands selected as the report type.
- Autonomous system (AS) Demand reports—To generate these reports, you must first enable the generation of AS demands by using the `set northstar analytics netflowd generate-as-demands` CLI command. Then, you must run the Demand Reports task, with the **Include AS Demands** toggle button enabled and with the required AS report types selected.

To run a demand reports task, see ["Add a Demand Reports Task" on page 985](#).

[Table 162 on page 948](#) explains the different types of demand reports. For details about each type of demand report, see ["View Demand Reports" on page 948](#).

Table 162: Types of Demand Reports

Report Type	Description
VPN Demands	Lists the amount of traffic per VPN, based on the selected time period.
Group Demands	Lists the amount of traffic per node group, based on the selected time period.
LSP Demands	Lists the amount of traffic per LSP, based on the selected time period.
AS Demands	Lists the amount of traffic per AS, based on the selected time period.

RELATED DOCUMENTATION

[NetFlow Collector Overview | 864](#)

[About the Demand Tab | 765](#)

View Demand Reports

Paragon Automation enables you to generate and view Demand reports that help you analyze network traffic. For information on the types of Demand reports and how to generate them, see "[Demand Reports Overview](#)" on page 947.

To view demand reports, select **Reports > Demand**.

You can perform the following actions on this page:

- View the following demand reports:
 - VPN Demand reports, see [Table 163 on page 949](#).
 - Group Demand reports, see [Table 164 on page 950](#).
 - Label-switched paths (LSP) Demand reports, see [Table 165 on page 950](#).
 - Autonomous system (AS) Demand reports, see [Table 166 on page 950](#).

- Download a report—Hover over the Download button on the top-right corner of the report and select either **CSV** or **JSON** to download the report as a comma-separated values file (CSV) or JavaScript Object Notation (JSON) file, respectively.
- Show or hide columns in a report—Click the vertical ellipsis icon at the top right corner of the report and select the columns that you want displayed in the report.

Only the columns that you selected are included in the report.
- View traffic in a time series graph—Select a demand and click **More > Show Chart**. The Time Series page appears, displaying a time series graph that shows the amount of traffic (in bps) for the time range for which you've generated the demand report. The x-axis indicates the time, while the y-axis indicates the amount of traffic.

Table 163 on page 949 explains the fields displayed in the VPN Demand reports.

NOTE: In addition to the fields described in the following tables, each of these reports displays the amount of traffic flow (in bps). Based on the period that you specified in the Demand Traffic Schedule field when adding the demand report task, you may see one or multiple columns displaying this data:

- Data pertaining to each hour, for the past 24 hours, if you've selected the **Demand Range for last 24 hours** option.
- Data pertaining to the past N days, if you've selected the **Demand Range for N days** option.
- Data pertaining to the selected time period, if you've selected the **Date Range** option.

Table 163: VPN Demand Reports

Field	Description
Routing Table	The name of the global routing table (default) or a routing-instance (like VRFs), where the ingress interface is configured. The ingress interface is an interface where sampling is performed.
Type	The type of demand data (LDP or IP) collected.
Ingress PE	The name of the provider edge (PE) device from which traffic flows.
Egress PE	The name of the PE device to which traffic flows.

Table 163: VPN Demand Reports (Continued)

Field	Description
Next Hop	The IP address of the BGP next hop.

[Table 164 on page 950](#) explains the fields displayed in the Group Demand reports.

Table 164: Group Demand Reports

Field	Description
Group	The name of the node group for which the report is generated.

[Table 165 on page 950](#) explains the fields displayed in the LSP Demand reports.

Table 165: LSP Demand Reports

Field	Description
LER: LSP	The name of the label-edge router (LER), which is the ingress node from which the LSP originates, and the LSP name.

[Table 166 on page 950](#) explains the fields displayed in the AS Demand reports.

NOTE: Based on the types of AS reports that you've selected when adding the demand reports task, you may see one or more of the following fields.

Table 166: AS Demand Reports

Field	Description
Ingress PE	Name of the PE device from which traffic flows.
Egress PE	Name of the PE device to which traffic flows.

Table 166: AS Demand Reports (Continued)

Field	Description
Ingress AS	Name of the AS from which traffic flows.
Egress AS	Name of the AS to which traffic flows.

RELATED DOCUMENTATION

| [Add a Demand Reports Task](#) | 985

11

PART

Administration

[Manage E-mail Templates](#) | 953

[Manage Audit Logs](#) | 957

[Configure External EMS](#) | 963

[Manage Task Scheduler](#) | 969

[Manage Security Settings](#) | 1001

[License Management](#) | 1005

Manage E-mail Templates

IN THIS CHAPTER

- [E-mail Templates Overview | 953](#)
- [About the E-mail Templates Page | 954](#)
- [Edit an E-mail Template | 955](#)

E-mail Templates Overview

Paragon Automation provides e-mail templates to communicate events, such as user account created, password changed, and so on, related to a user account. The e-mail templates are mapped to specific events and when the specific event occurs, the corresponding e-mail is sent from Paragon Automation to the user. For example, when a user account is created, the activation mail is sent out to the user informing about the user's username, a link to set the password and the link to access the Paragon Automation portal.

You must first configure SMTP on Paragon Automation for the e-mails to be sent to the user. For information about configuring SMTP settings, see "[Configure SMTP Settings](#)" on page 50.

[Table 167 on page 953](#) describes the default e-mail templates provided by Paragon Automation.

Table 167: Default E-mail Templates

Template Name	Description
Account Locked	A mail is sent using this template when a user attempts to log in by using an incorrect password for more than five times.
Activation User Link	A mail is sent using this template when you resend the activation link to a user.
Changed Password	A mail is sent using this template when a user changes their password.

Table 167: Default E-mail Templates (Continued)

Template Name	Description
Forgot Account Password	A mail is sent using this template when a user clicks the Forgot Password link in the Paragon Automation Login page. The e-mail has the link the reset the password.
Invite User	A mail is sent using this template when a user who is already added to a tenant adds a user to another tenant. NOTE: This template is not applicable for this release.
Rest Your Account Password	A mail is sent using this template for a user to reset their Paragon Automation account password. The account password must be changed once every three months.
New Account Created	A mail is sent using this template when a new user account is created. The e-mail contains the link to set password, the username to log in to the Paragon Automation GUI and the URL of the Paragon Automation GUI.
Activate User	A mail is sent to you, when you need to unlock a user or activate a user whose invite has expired. The mail contains the link to activate the user.

RELATED DOCUMENTATION

| [About the E-mail Templates Page | 954](#)

About the E-mail Templates Page**IN THIS SECTION**

- [Tasks You Can Perform | 955](#)
- [Field Descriptions | 955](#)

To access this page, click **Administration > Authentication > Email Templates**.

Use this page to view and edit predefined e-mail templates. An e-mail is sent to the users when a specific event occurs.

Tasks You Can Perform

You can perform the following tasks on this page:

- View e-mail templates. To view details of a specific e-mail template, select the e-mail template and click **More > Detail**. Alternatively, hover over the e-mail template name and click the **Details** icon that appears. The *Details for Template-Name* page is displayed.
- Edit e-mail templates. For more information see ["Edit an E-mail Template" on page 955](#).

Field Descriptions

[Table 168 on page 955](#) displays the fields on the Email Templates page.

Table 168: Fields on the Email Templates Page

Field	Description
Template Name	The name of the e-mail template.
Modified By	The user who last modified the e-mail template.
Last Updated	The date and time when the e-mail template was last updated.
Content	The content of the e-mail template.

RELATED DOCUMENTATION

| [E-mail Templates Overview | 953](#)

Edit an E-mail Template

If you are an administrator, you can edit the e-mail templates. To edit an existing e-mail template:

1. Select **Authentication > Email Templates**.

The Email Templates page is displayed.

2. Select the e-mail template that you want to edit and click the **edit (pencil)** icon.

The Edit Email Template page appears.

3. Edit the content of the e-mail.

The e-mail template is in the Jinja format. You cannot edit the name of the e-mail template.

NOTE: If you want to view the default content at this stage, you can click **Restore Default Content**. However, the edits that you have made will not be saved.

4. Click **OK** to save the changes that you have made to the default content.

You have edited the content in the e-mail template.

RELATED DOCUMENTATION

| [E-mail Templates Overview](#) | 953

Manage Audit Logs

IN THIS CHAPTER

- [Audit Logs Overview | 957](#)
- [About the Audit Logs Page | 959](#)
- [Filter Audit Logs | 960](#)
- [Export Audit Logs | 961](#)

Audit Logs Overview

IN THIS SECTION

- [Paragon Insights Commands and Audit Logs | 958](#)

An audit log is a record of a sequence of user activities that affect a specific operation. Audit logs are useful for tracing events and for maintaining a record of the user's activities. Network operators can view or filter audit logs to determine which user performed what action at a given time. The most recent audit log appears first.

Audit logs contain information about the tasks initiated by a user from the Paragon Automation GUI or APIs. Audit log entries include:

- The name of the user who initiated the task
- Source IP address of the client where the task was initiated
- Operation that triggered the audit log
- Status of the task
- Short description of the task

- Related job details. Click the **Job ID** to view the job details on the **Job Status** page.

The job details include the job name, start and end time, job state (failed or success), job ID, and the related task details.

- Date and time of the task execution

Audit logs from microservices, such as healthbot command and config server are collected in the audit log client library, and are stored in the Postgres SQL database. The audit log client library is installed with microservices during Paragon Automation installation.

Paragon Pathfinder operations, for example, adding, modifying, or deleting network elements (such as nodes or links) are logged here.

NOTE: Device-driven tasks (that is, tasks not initiated by the user) are not recorded in audit logs.

Paragon Insights Commands and Audit Logs

Audit logs are generated when you run the following commands from the CLI:

- add-node
- remove-node
- modify-uda-engine
- modify-udf-engine
- modify-workflow-engine
- remove-plugin
- load-plugin

The audit log generated includes information on which user initiated the job, the node name, and status of the job. However, you must enter credentials (username and password) to run the commands. If you have already set a username (HB_USERNAME) and password (HB_PASSWORD), you are not prompted to enter credentials when you run a command.

RELATED DOCUMENTATION

| [About the Audit Logs Page](#) | 959

About the Audit Logs Page

IN THIS SECTION

- [Tasks You Can Perform | 959](#)

To access this page, select **Administration > Audit Logs**.

Use the **Audit Logs** page to view the tasks that you have initiated either by using the Paragon Automation GUI or APIs. You can export audit logs as a comma-separated values (CSV) file.

Tasks You Can Perform

You can perform the following tasks from this page:

- View the details of various user-initiated tasks by selecting particular audit log record then clicking on **More > Details**. You can also mouse over the audit log and click on the **Detailed View** icon. See [Viewing the Details of an Audit Log](#).
- Export audit logs as a CSV file. See ["Export Audit Logs" on page 961](#).
- Sort and filter audit logs:

NOTE: Sorting and filtering is applicable only to some fields.

- Click a column name to sort the audit logs based on the column name.
- Click the **filter** icon and select whether you want to show or hide column filters or apply a quick filter. For example, you can use audit log filtering to track user accounts that were added on a specific date, track configuration changes across a particular type of device, view services that were provisioned on specific devices. See ["Filter Audit Logs" on page 960](#).
- Show or hide columns—Click the **Show Hide Columns** icon at the top right corner of the page and select the columns that you want to display on the **Audit Logs** page.

[Table 169 on page 960](#) provides description of the fields on the Audit Logs page.

Table 169: Fields on the Audit Logs Page

Field	Description
Username	Displays the username of the user who initiated the task.
Source IP	Displays the IP address of the client from which the user initiated the task. For tasks that do not have an associated Source IP address, this field is blank.
Object Name	Displays the name of the object on which the task was initiated. An object can be a tenant, site, device, device image, template, and so on.
Operation	Displays the name of the task that triggered the audit log.
Description	Displays details about the task.
Status	Displays the status of the task that triggered the audit log: <ul style="list-style-type: none"> • Success—Job or task was completed successfully. • Failure—Job or task failed and was terminated.
Logged time	Time at which audit log is recorded.
Job ID	Job ID contains a uniquely identified job which is associated with the audit log activity. The Job ID value is a hyperlink, clicking it will open the Job overlay page which displays complete information about that particular job

Filter Audit Logs

You can filter audit logs from the **Administration > Audit Logs** page of the Paragon Automation GUI. You can apply filters to audit logs before you export audit logs in a CSV file. For more information, see ["Export Audit Logs" on page 961](#).

Follow these steps to filter audit logs:

1. Select **Administration > Audit Logs**.

The **Audit Logs** page appears displaying the audit logs.

2. Click **Filter**.
The **Add Criteria** pop-up appears.
3. Enter the following information in the **Add Criteria** pop-up.
 - a. Select the column you want to apply the filter to from the **Field** drop-down list.
 - b. Select the condition you want to apply the column from the **Condition** drop-down list.
 - c. Enter the start or finish time in the **Value** text box, that you want to apply to the filter.
4. Click **Add** to apply the filter.
(Optional) You can export audit logs as a CSV file after you apply a filter. For more information, see ["Export Audit Logs" on page 961](#).

Export Audit Logs

You can export audit logs as comma-separated values (CSV) file that can be opened or edited using an application such as Microsoft Excel.

You can view and analyze the exported audit logs, as needed.

To export the audit logs:

1. Select **Administration > Audit Logs**. The **Audit Logs** page appears displaying the audit logs.
2. Click **Export Logs > CSV** to download audit logs in as a CSV file.
In the warning message that appears, do one of the following:
 - Click **Continue** to export all audit logs with all fields as a CSV file.
 - Click **Filter** to filter audit log fields before export. For more information, see ["Filter Audit Logs" on page 960](#).

After you have applied a filter, click **Export Logs > CSV** again to start the export.

NOTE: You can export audit logs for a maximum of 30 days prior to the current date and time. For example, if the current date is May 31, 2018, you can export the audit logs starting from May 1, 2018.

Depending on the settings of the browser that you are using, the CSV file containing the audit logs for the specified time period is either downloaded directly, or you are asked to open or save the file.

RELATED DOCUMENTATION

| [About the Audit Logs Page](#) | **959**

Configure External EMS

IN THIS CHAPTER

- [About the External EMS Page | 963](#)
- [Add an External EMS | 965](#)
- [Edit and Delete an External EMS | 967](#)

About the External EMS Page

IN THIS SECTION

- [Tasks You Can Perform | 964](#)
- [Field Descriptions | 965](#)

External EMS refers to a software such as Anuta ATOM, Contrail Service Orchestration, Junos Space Security Director, or any other open source software with EMS capabilities, that can be integrated with Paragon Automation. When you integrate an External EMS with Paragon Automation, the devices managed by the External EMS are added to Paragon Automation and you can manage those devices by using Paragon Automation.

You can use this feature when you already have an EMS application managing your devices and you plan to use Paragon Automation for managing devices going forward.

NOTE: Anuta ATOM is the only external EMS supported in Paragon Automation.

To access this page, click **Administration > External EMS**.

Tasks You Can Perform

You can perform the following tasks from this page:

- Add an External EMS application; see ["Add an External EMS" on page 965](#).
- Edit and delete an External EMS application from Paragon Automation; see ["Edit and Delete an External EMS" on page 967](#).
- Filter Entries—Filter the table entries by adding new filtering criteria.

Hover over the **Filter** (funnel) icon and select **Add Filter**. On the **Add Criteria** page, select the filtering criteria from the **Field** and **Condition** list, and enter the text to be compared in the **Value** field. Click **Add**.

The filtered table entries are listed and the filter criteria name is displayed above the table column names.

You can click the cross (X) icon (next to the filter name) to remove the filtering criteria.

NOTE:

- You can add multiple filtering criteria. Once you add the multiple filtering criteria, you can select the **And** condition to display the entries matching all the filtering criteria or select the **Or** condition to display the entries matching any one of the filtering criteria.
- Quick filter: Save the filtering criteria as quick filters. Once you have added all the filtering criteria, you can save a particular criteria or multiple criteria for future use by clicking **Save**.

On the **Save Filter** window, 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 a specific column in the table.

Hover over the **More Options** (vertical ellipsis) > **Show/Hide Columns** and select the *Column-Name* check box of the columns you want to display in the table.

- Reset Preference—Reset the displayed columns to the default set of columns for each tab in the table.

Hover over the **More Options** (vertical ellipsis) and select **Reset Preference**.

- **Sort Entries**—Click the column name to highlight the up and down arrows next to the column name. Sort the table entries in ascending or descending order of that column by clicking the up or down arrow respectively.

Field Descriptions

Table 170 on page 965 lists the fields on the External EMS page.

Table 170: Fields on the External EMS Page

Field	Description
Type	The External EMS application.
Server Address	The IPv4 address to access the External EMS application.
Username	The username to log in to the External EMS application.
Parameters	The parameters such as client ID and client secret configured for connecting with the External EMS application.

RELATED DOCUMENTATION

[Integration with Anuta ATOM Overview | 90](#)

[About the Identity Providers Page | 81](#)

Add an External EMS

Only an administrator or a user with privileges to add an External EMS application can add an External EMS application to Paragon Automation.

Before you add an External EMS to Paragon Automation, ensure that you have an account created in the EMS application. To add the external EMS application to Paragon Automation, you need to provide the credentials of your account in the external EMS application.

To add an external EMS to Paragon Automation:

1. Click **Administration >External EMS** in the left navigation menu.

The External EMS page appears.

2. Enter values referring to [Table 171 on page 966](#).
3. Click **OK**.

A message indicating whether the External EMS application was successfully added or not is displayed. The External EMS application is listed on the External EMS page if the operation is successful.

[Table 171 on page 966](#) lists the fields on the Add External EMS page.

Table 171: Fields on the Add External EMS Page

Field	Description
Type	Select the External EMS application that you want to integrate with Paragon Automation from the list. ATOM is the only supported external EMS in Paragon Automation.
Server Address	Enter the IPv4 address to access the External EMS application.
Username	Enter the username to log in to the External EMS application.
Password	Enter the password to log in to the External EMS application.
Parameters	<p>Add parameters, such as client ID and client secret, for connecting with the External EMS application.</p> <p>To add the parameters, click the Add icon (+) provided in the Parameters field and enter the parameter and corresponding value in the text fields that appear. Add more parameters by clicking the Add (+) icon.</p>

RELATED DOCUMENTATION

[Register Paragon Automation with Keycloak | 93](#)

[Add Identity Providers | 84](#)

Edit and Delete an External EMS

IN THIS SECTION

- [Edit an External EMS | 967](#)
- [Delete an External EMS | 967](#)

Only an administrator or a user with privileges to edit or delete an External EMS application can edit or delete an External EMS application.

Edit an External EMS

You can edit the following parameters of an external EMS:

- IPv4 address to access the external EMS application.
- Username and password to log in to the external EMS application.
- Any parameters that are configured for connecting with the external EMS.

To edit an External EMS application:

1. Click **Administration > External EMS** in the left navigation menu.

The External EMS page appears.

2. Select the EMS that you want to edit and click the **Edit** (Pencil) icon .

The Edit External EMS page appears.

3. Edit the parameters by referring to [Table 171 on page 966](#) and click **OK**.

A message appears indicating that the external EMS application is edited successfully and the updates are displayed on the External EMS page.

Delete an External EMS

Only an administrator or a user with privileges to delete an External EMS application can delete an External EMS application.

To delete an EMS application that is integrated with Paragon Automation:

1. Click **Administration > External EMS** in the left navigation menu.

The External EMS page appears.

2. Select the External EMS application that you want to delete and click the **Delete** icon (Trashcan).

A confirmation message appears asking you to confirm whether you want to delete the application.

3. Click **OK.**

A message appears indicating that the external EMS application is deleted successfully and the External EMS application is no longer listed on the External EMS page.

SEE ALSO

[About the External EMS Page | 963](#)

[Register Paragon Automation with Keycloak | 93](#)

Manage Task Scheduler

IN THIS CHAPTER

- [About the Task Scheduler Page | 969](#)
- [Add a Bandwidth Sizing Task | 972](#)
- [Add a Container Normalization Task | 975](#)
- [Add a Device Collection Task | 978](#)
- [Add a Demand Aging Task | 983](#)
- [Add a Demand Reports Task | 985](#)
- [Add a Network Archive Task | 990](#)
- [Add a Network Maintenance Task | 993](#)
- [Add a Network Cleanup Task | 997](#)
- [Edit and Delete Tasks | 999](#)

About the Task Scheduler Page

IN THIS SECTION

- [Tasks You Can Perform | 970](#)
- [Field Descriptions | 971](#)

To access this page, select **Administration > Task Scheduler**.

Use the Task Scheduler page to view and manage custom (user-created) tasks in Paragon Pathfinder. You can only view the system tasks that the Path Computation Element (PCE) launches to run scripts; you cannot add, modify, or delete the system tasks.

The system tasks are:

- **CollectionCleanup:** Purges old raw and aggregated analytics data.
- **ESRollup:** Aggregates the collected data from the previous hour.

Tasks You Can Perform

You can perform the following tasks from this page:

- View the list of scheduled and completed tasks, and their details—See [Table 172 on page 971](#) for information.
- Add tasks—Currently, you can add the following tasks:

Task Group - Bandwidth Management:

- **Bandwidth sizing task**—See ["Add a Bandwidth Sizing Task" on page 972](#).
- **Container Normalization**—See ["Add a Container Normalization Task" on page 975](#).

Task Group - Collection Tasks:

- **Device Collection**—See ["Add a Device Collection Task" on page 978](#).

Task Group - Report Tasks:

- **Demand Aging**—See ["Add a Demand Aging Task" on page 983](#).
- **Demand Reports**—See ["Add a Demand Reports Task" on page 985](#).

Task Group - Utility Tasks:

- **Network Archive**—See ["Add a Network Archive Task" on page 990](#).
- **Network Maintenance**—See ["Add a Network Maintenance Task" on page 993](#).
- **Network Cleanup**—See ["Add a Network Cleanup Task" on page 997](#).
- **Modify the parameters configured for a task**—See ["Edit and Delete Tasks" on page 999](#).
- **Delete existing tasks**—See ["Edit and Delete Tasks" on page 999](#).
- **Show or hide columns displayed on the page**—Click the Custom (vertical ellipsis) icon and select **Show/Hide Columns**. Then, select the check boxes corresponding to the columns that you want to view.

The columns that you selected are displayed on the page.

Field Descriptions

Table 172: Fields on the Task Scheduler Page

Field	Description
ID	ID generated for the task.
Type	Type of the task.
Name	Name of the task. For system tasks, the name can be CollectionCleanup or ESRollup Task .
Created	Date (in YYYY:MM:DD format) and time (HH:MM:SS 24-hour format) at which the task was created.
Frequency	Periodicity of the recurrence (Immediately, Once, Minutes, Hourly, Daily, Weekly, Monthly, or Yearly).
Repeats	The number of minutes, hours, days, weeks, months, or years after which the task is scheduled to recur. This value is displayed only if the recurrence is scheduled. Displays N/A if the task is scheduled to be executed immediately. Displays Never if the task is not scheduled to recur.
Starts	Date (in YYYY:MM:DD format) and time (HH:MM:SS 24-hour format) at which the task starts.
Ends	Date (in YYYY:MM:DD format) and time (HH:MM:SS 24-hour format) at which the task completes. This value is displayed only if you've specified the end time for the task. Displays N/A if the task is scheduled to be executed immediately. Displays Never if the task is scheduled to never end.

Table 172: Fields on the Task Scheduler Page (Continued)

Field	Description
Last Executed	Date (in YYYY:MM:DD format) and time (HH:MM:SS 24-hour format) at which the task was last executed.
Status	Current status of the task (Running, Scheduled, or Completed).
System Task	Indicates whether the task is a system task (True) or user-created task (False).

Add a Bandwidth Sizing Task

The bandwidth sizing task periodically sends a new planned bandwidth for bandwidth sizing-enabled tunnels (also known as label-switched paths or LSPs) to the Path Computation Server (PCS). The PCS determines whether it needs to provision the new planned bandwidth with a path that satisfies the new bandwidth requirement.

To add a bandwidth sizing task:

1. Select **Administration > Task Scheduler**.

The Task Scheduler page appears.

2. Click the + (Add) icon.

The Add Task page (wizard) appears.

NOTE: You can have only one bandwidth sizing task per Paragon Automation server. If you attempt to add a second, the system will prompt you to approve overwriting the first one.

3. Configure the fields in each step of the wizard according to the guidelines provided in [Table 173 on page 973](#).

Click **Next** to go to the next step.

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

4. After you complete each step in the wizard, click **Finish**.

A confirmation message appears on the top of the page, indicating that the task is added successfully.

The details of this task are displayed on the Task Scheduler page.

Table 173: Fields on the Add Task Wizard (Bandwidth Sizing)

Field	Description
<i>Add New Task</i>	
Task Name	Specify a unique name for the task. The name can contain only alphanumeric characters and some special characters (greater than (>), less than (<), colon (:), underscore (_), and hyphen (-)).
Task Group	From the list, select Bandwidth Management Tasks .
Task Type	From the list, select Bandwidth Sizing .
<i>Bandwidth Sizing Task</i>	
Aggregation Statistic	<p>From the list, select an aggregation statistic option.</p> <p>The aggregation statistic works together with the task execution recurrence interval (the period of bandwidth adjustment) that you specify in the Schedule step in this wizard. The Path Computation Element (PCE) aggregates the LSP traffic for the interval based on the aggregation statistic you select, and uses that information to calculate the new planned bandwidth. The available options are:</p> <ul style="list-style-type: none"> 80th, 90th, 95th, 99th Percentile (X percentile)—Aggregation is based on the selected percentile. <p>The 'X' percentile is the value at which 'X' percent of all the samples taken in the previous sampling period lie at or below the calculated value.</p> <p>For bandwidth sizing, the newly-calculated bandwidth value is taken as the 'X' percentile of the samples in the immediately-preceding bandwidth sizing interval.</p> <ul style="list-style-type: none"> Average—For each interval, the samples within that interval are averaged. If there are N samples for a particular interval, the result is the sum of all the sample values divided by N. Max—For each interval, the maximum of the sample values within that interval is used.

Table 173: Fields on the Add Task Wizard (Bandwidth Sizing) (Continued)

Field	Description
<i>Schedule</i>	
Startup Options	<p>Schedule—Select Activate Later for the task to start at a specific date and time.</p> <p>In the field that appears, click the calendar icon to select the date and time at which you want the task to start. Then, click OK.</p> <p>For the task to start at the current date and time, click Now at the bottom of the calendar.</p> <p>NOTE:</p> <ul style="list-style-type: none"> • You can schedule the task to repeat at a specific interval, from a minimum of 15 minutes to a maximum of one day. The default interval is one hour. • There is no per-LSP interval. The interval configured here applies to all LSPs for which bandwidth sizing is enabled. • LSPs for which traffic statistics are not available for a specified duration will not be resized if the bandwidth sizing task is scheduled to consider the statistics collected in the specified duration. This is because the task cannot differentiate between no traffic and zero bandwidth traffic.
Recurrence Options	<ul style="list-style-type: none"> • Repeats—Specify the frequency (in Minutes, Hours, or Days) at which the task recurs. • Every—Specify the periodicity of the recurrence. • Ends—Select one of the following options: <ul style="list-style-type: none"> • Never—To configure the task to recur at the specified interval. • On—To select a date and time at which you want the task to end. In the field that appears, click the calendar icon and select the date and time. Then, click OK. • Now—To end a task at the current date and time, click Now at the bottom of the calendar.

RELATED DOCUMENTATION

[About the Task Scheduler Page | 969](#)

[Edit and Delete Tasks | 999](#)

Add a Container Normalization Task

Add a container normalization task for periodic container label-switched path (LSP) normalization.

The container normalization task computes aggregated bandwidth for each container LSP and sends it to the Path Computation Server (PCS). The PCS determines whether it needs to add or remove sub-LSPs belonging to the container LSP, based on the container's new aggregated bandwidth.

To add a container normalization task:

1. Select **Administration > Task Scheduler**.

The Task Scheduler page appears.

2. Click the + (Add) icon.

The Add Task page (wizard) appears.

NOTE: You can have only one container normalization task per Paragon Automation server. If you attempt to add a second, the system will prompt you to approve overwriting the first one.

3. Configure the fields in each step of the wizard according to the guidelines provided in [Table 174 on page 975](#).

Click **Next** to go to the next step.

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

4. After you complete each step in the wizard, click **Finish**.

A confirmation message appears on top of the page, indicating that the task was added successfully.

The details of this task are displayed on the Task Scheduler page.

Table 174: Fields on the Add Task Wizard (Container Normalization)

Field	Description
<i>Add New Task</i>	
Task Name	Specify a unique name for the task. The name can contain only alphanumeric characters and some special characters (greater than (>), less than (<), colon (:), underscore (_), and hyphen (-)).

Table 174: Fields on the Add Task Wizard (Container Normalization) (Continued)

Field	Description
Task Group	From the list, select Bandwidth Management Tasks .
Task Type	From the list, select Container Normalization .

Container Normalization Task

Aggregation Statistic	<p>From the list, select an aggregation statistic option.</p> <p>The aggregation statistic works together with the task execution recurrence interval (the period of bandwidth adjustment) that you specify in the Schedule step in this wizard. The Path Computation Element (PCE) aggregates the LSP traffic for the interval based on the aggregation statistic you select, and uses that information to calculate the new aggregated bandwidth. The available options are:</p> <ul style="list-style-type: none"> • 80th, 90th, 95th, 99th Percentile (X percentile)—Aggregation is based on the selected percentile. <p>The 'X' percentile is the value at which 'X' percent of all the samples taken in the previous sampling period lie at or below the calculated value.</p> <ul style="list-style-type: none"> • Average—For each interval, the samples within that interval are averaged. If there are N samples for a particular interval, the result is the sum of all the sample values divided by N. • Max—For each interval, the maximum of the sample values within that interval is used.
-----------------------	--

Schedule

Table 174: Fields on the Add Task Wizard (Container Normalization) (Continued)

Field	Description
Startup Options	<p>Schedule—Select Activate Later for the task to start at a specific date and time.</p> <p>In the field that appears, click the calendar icon to select the date and time at which you want the task to start. Then, click OK.</p> <p>For the task to start at the current date and time, click Now at the bottom of the calendar.</p> <p>NOTE:</p> <ul style="list-style-type: none"> You can schedule the task to repeat at a specific interval, from a minimum of 15 minutes to a maximum of one day. The default interval is one hour. Container LSPs for which traffic statistics are not available for a specified duration will not be normalized if the container normalization task is scheduled to consider the statistics collected in the specified duration. This is because the task cannot differentiate between no traffic and zero bandwidth traffic.
Recurrence Options	<ul style="list-style-type: none"> Repeats—Specify the frequency (in Minutes, Hours, or Days) at which the task recurs. Every—Specify the periodicity of the recurrence. Ends—Select one of the following options: <ul style="list-style-type: none"> Never—To configure the task to recur at the specified interval. On—To select a date and time at which you want the task to end. In the field that appears, click the calendar icon and select the date and time. Then, click OK. Now—To end a task at the current date and time, click Now at the bottom of the calendar.

RELATED DOCUMENTATION

[About the Task Scheduler Page | 969](#)

[Edit and Delete Tasks | 999](#)

Add a Device Collection Task

The Analytics features in Pathfinder require that the Path Computation Server (PCS) periodically connect to the network in order to obtain the configuration of network devices. The PCS uses this information to correlate IP addresses, interfaces, and devices. In addition, the PCS also obtains additional information about the devices. For example, VPN routing instances configured on the devices can be discovered and parsed as a result of this task.

Adding of devices (**Configuration > Devices**) to Paragon Automation is a prerequisite for successfully running device collection tasks.

NOTE: For topologies that include logical nodes, periodic device collection is necessary as there are no real-time PCEP-based updates for logical devices.

To add a device collection task:

1. Select **Administration > Task Scheduler.**

The Task Scheduler page appears.

2. Click the + (Add) icon.

The Add Task page (wizard) appears.

3. Configure the fields in each step of the wizard according to the guidelines provided in [Table 175 on page 978](#).

Click **Next** to go to the next step.

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

4. After you complete each step in the wizard, click **Finish.**

A confirmation message appears on top of the page, indicating that the task was added successfully.

The details of this task are displayed on the Task Scheduler page. The device collection data is sent to the PCS for routing and is reflected in the Topology view.

Table 175: Fields on the Add Task Wizard (Device Collection)

Field	Description
-------	-------------

Add New Task

Table 175: Fields on the Add Task Wizard (Device Collection) (Continued)

Field	Description
Task Name	Specify a unique name for the task. The name can contain only alphanumeric characters and some special characters (greater than (>), less than (<), colon (:), underscore (_), and hyphen (-)).
Task Group	From the list, select Collection Tasks .
Task Type	From the list, select Device Collection .
<i>Device Collection Task</i>	
Task Options > Devices to be collected	By default, all devices are selected to be included in the collection task (that is, the All Devices toggle button is disabled). Click the toggle button to enable including only selective devices in the device collection task. If you enable this toggle button, a list of all the devices available to be included in the collection task is displayed. Select the check boxes corresponding to the devices you want to include.
Task Options > Other Options	
Archive BGP Files	If you select this check box, the memory consumption of large BGP files in Pathfinder is reduced by compressing data collection files that result from the device collection task. Default: Selected (Yes)
Archive Raw Data	If you select this check box, raw data is archived in Elasticsearch. Default: Selected (Yes)
Store Collection for Planner	If you select this check box, raw and spec data are added to the database, making it available for import into Paragon Planner as a network. This data includes the unparsed collected data from the devices and the device collection task status. The unparsed collected data contains the output of multiple show commands run on the devices. When you select this check box, the Network Description field becomes available.
Network Description	(Optional) Add a meaningful description for your reference.

Table 175: Fields on the Add Task Wizard (Device Collection) (Continued)

Field	Description
Parse Collection	<p>If you select this check box, Pathfinder reads the content of the configuration files and updates the network model accordingly. If you don't select this option, the configuration files are collected on the server, but not used in the model.</p> <p>Default: Selected (Yes)</p>
Use Management IP	<p>The behavior of this option depends on whether the management IP is configured or not:</p> <p>Default: Selected (Yes)</p> <p>If you select Use management IP and:</p> <ul style="list-style-type: none"> • The management IP is configured—Only the management IP is tried for collection, whether it is reachable or not. • The management IP is not configured—The IP address is tried for collection. <p>If you don't select Use management IP and:</p> <ul style="list-style-type: none"> • The management IP is configured—The management IP is tried first for collection. If the management IP is configured, but is not accessible, the IP address is tried for collection. • The management IP is not configured—The IP address is tried for collection.
Collection Options	
Configuration Collection	<p>Click the toggle button to enable or disable (default) configuration collection.</p> <p>If you enable this toggle button, router configuration is automatically collected. In addition, the Interface, Tunnel Path, and Transit Tunnel data are collected by default. You can select or clear one or more check boxes corresponding to the data to be collected or processed.</p> <p>NOTE: We recommend that you collect router configuration, tunnel path, and tunnel transit data when running the device collection task so that the PCE can update the tunnel status and details based on the latest collection.</p> <p>Each of the options results in the collection task capturing the results of various show commands. Table 176 on page 982 lists the show command output captured for each option.</p>

Table 175: Fields on the Add Task Wizard (Device Collection) (Continued)

Field	Description
CLI Collection	<p>Equipment CLI—If you select this check box, equipment CLI data is collected as part of this task. The Process Equipment CLI option in the Network Archive task parses the Equipment CLI data that is collected during device collection. The Inventory Report that is generated is available in both Paragon Pathfinder and Paragon Planner.</p> <p>To view Hardware Inventory in Paragon Planner, you must run device collection with the Equipment CLI collection option (collects the inventory data) and the Network Archive collection with the Process Equipment CLI option (processes the inventory data).</p>
<i>Schedule</i>	
Startup Options	<p>Schedule—Select one of the following options to schedule the task:</p> <ul style="list-style-type: none"> • Activate Now: The task starts at the current date and time. • Activate Later: The task starts at a later date and time. <p>In the field that appears, click the calendar icon to select the date and time at which you want the task to start. Then, click OK.</p> <p>Instead of using the Activate Now option, you can also click Now at the bottom of the calendar for the task to start at the current date and time.</p> <p>NOTE: You can choose to run the collection only once, or to repeat it at specific intervals. The default interval is 15 minutes.</p>

Table 175: Fields on the Add Task Wizard (Device Collection) (Continued)

Field	Description
Recurrence Options	<ul style="list-style-type: none"> • Repeats—Specify the frequency (in Minutes, Hours, Days, Weekly, Monthly, Yearly, or Never) at which the task recurs. • Every—Specify the periodicity of the recurrence. • Ends—Select one of the following options: <ul style="list-style-type: none"> • Never—To configure the task to recur at the specified interval. • On—To select a date and time at which you want the task to end. In the field that appears, click the calendar icon and select the date and time. Then, click OK. • Now—To end a task at the current date and time, click Now at the bottom of the calendar.

Table 176 on page 982 lists the show command output captured for each Collection Option.

Table 176: Show Command Output Captured by Device Collection Options

Data Type	For Juniper Devices	For IOS-XR Devices
Configuration	show configuration display inheritance brief no-more	show running
Interface	show configuration system host-name display inheritance brief show interfaces no-more	show running include hostname show interfaces show ipv4 interface
Tunnel Path	show configuration system host-name display inheritance brief show mpls lsp statistics ingress extensive logical-router all no-more	show running include hostname show mpls traffic-eng tunnels detail role head

Table 176: Show Command Output Captured by Device Collection Options (Continued)

Data Type	For Juniper Devices	For IOS-XR Devices
Transit Tunnel	show configuration system host-name display inheritance brief show rsvp session ingress detail logical- router all no-more show rsvp session transit detail logical- router all no-more	show running include hostname show mpls traffic-eng tunnels backup
Equipment CLI	show configuration system host-name display inheritance brief show version no-more show chassis hardware no-more show chassis fpc no-more show chassis hardware models no-more	show version show diag show env all admin show inventory show inventory raw

RELATED DOCUMENTATION

[About the Task Scheduler Page | 969](#)

[Edit and Delete Tasks | 999](#)

Add a Demand Aging Task

You can add a Demand Aging task to automatically delete inactive demands (according to the maximum age you specify) from the Demand tab in the network information table.

To add a demand aging task:

1. Select **Administration > Task Scheduler**.

The Task Scheduler page appears.

2. Click the + (Add) icon.

The Add Task page (wizard) appears.

3. Configure the fields in each step of the wizard according to the guidelines provided in [Table 177 on page 984](#).

Click **Next** to go to the next step.

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

4. After you complete each step in the wizard, click **Finish**.

A confirmation message appears on top of the page, indicating that the task is added successfully.

The details of this task are displayed on the Task Scheduler page.

When this task is executed, demands that are no longer active are automatically deleted from the Demand tab in the network information table, based on the maximum age you specify.

Table 177: Fields on the Add Task Wizard (Demand Aging)

Field	Description
<i>Add New Task</i>	
Task Name	Specify a unique name for the task. The name can contain only alphanumeric characters and some special characters (greater than (>), less than (<), colon (:), underscore (_), and hyphen (-)).
Task Group	From the list, select the task group to which the task belongs. In this case, select Report Tasks .
Task Type	From the list, select the type of task that you want to add. In this case, select Demand Aging .
<i>Demand Aging Task</i>	
Demand Aging Options	<ul style="list-style-type: none"> • Maximum Demand Age—Specify the maximum time (integer value) for the demands to be deleted. For example, if you specify the maximum age as ten minutes, the task deletes all demands that have been inactive for ten minutes or more. • Units—From the list, select the units (Days, Hours, or Minutes) for the maximum demand age.

Table 177: Fields on the Add Task Wizard (Demand Aging) (Continued)

Field	Description
<i>Schedule</i>	
Startup Options	<p>Schedule—Select one of the following options to schedule the task:</p> <ul style="list-style-type: none"> • Activate Now: The task starts at the current date and time. • Activate Later: The task starts at a later date and time. <p>In the field that appears, click the calendar icon to select the date and time at which you want the task to start. Then, click OK.</p> <p>Instead of using the Activate Now option, you can also click Now at the bottom of the calendar for the task to start at the current date and time.</p>
Recurrence Options	<ul style="list-style-type: none"> • Repeats—Specify the frequency (in Minutes, Hours, Days, Weekly, Monthly, Yearly, or Never) at which the task recurs. • Every—Specify the periodicity of the recurrence. • Ends—Select one of the following options: <ul style="list-style-type: none"> • Never—To configure the task to recur at the specified interval. • On—To select a date and time at which you want the task to end. In the field that appears, click the calendar icon and select the date and time. Then, click OK. • Now—To end a task at the current date and time, click Now at the bottom of the calendar.

RELATED DOCUMENTATION

| [About the Demand Tab](#) | 765

Add a Demand Reports Task

You can add a demand reports task to generate demand reports that are listed in the Demand page (**Reports > Demand**).

To add a demand reports task:

1. Select **Administration > Task Scheduler**.
The Task Scheduler page appears.
2. Click the + (Add) icon.
The Add Task page (wizard) appears.
3. Configure the fields in each step of the wizard according to the guidelines provided in [Table 178 on page 986](#).
Click **Next** to go to the next step.

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

4. After you complete each step in the wizard, click **Finish**.
A confirmation message appears on top of the page, indicating that the task is added successfully.
The details of this task are displayed on the Task Scheduler page.

When this task is executed, one or more demand reports are generated, based on the Report Types you select. You can view these reports in the Demand page (**Reports > Demand**).

Table 178: Fields on the Add Task Wizard (Demand Reports)

Field	Description
<i>Add New Task</i>	
Task Name	Specify a unique name for the task. The name can contain only alphanumeric characters and some special characters (greater than (>), less than (<), colon (:), underscore (_), and hyphen (-)).
Task Group	From the list, select the task group to which the task belongs. In this case, select Report Tasks .
Task Type	From the list, select the type of task that you want to add. In this case, select Demand Reports .
<i>Demand Reports Task</i>	

Table 178: Fields on the Add Task Wizard (Demand Reports) (Continued)

Field	Description
Report Types	<ul style="list-style-type: none"> • Include Demands—From the list, select one or more types of reports to be generated. The VPN Demands and Group Demands are selected by default. You can also select LSP Demands. • Include AS Demands—To generate autonomous system (AS) demand reports, you must first enable the generation of AS demands by using the set northstar analytics netflowd generate-as-demands CLI command. Then, click to enable this toggle button to generate AS demand reports. This toggle button is disabled by default. If you enable this toggle button, the AS Report Types list appears. You can select one or more types of AS reports to be generated. The other available options are: <ul style="list-style-type: none"> • Ingress AS, Egress AS (default) • Ingress PE, Ingress AS (default) • Ingress PE, Ingress AS, Egress AS • Egress PE, Ingress AS, Egress AS • Ingress PE, Egress AS • Egress PE, Ingress AS • Egress PE, Egress AS • Ingress AS • Egress AS • Ingress PE, Ingress AS, Egress PE, Egress AS

Table 178: Fields on the Add Task Wizard (Demand Reports) (Continued)

Field	Description
Report Options > Demand Traffic Schedule	<p>Select one of the following options to specify the period for which you want the demand traffic data to be collected:</p> <ul style="list-style-type: none"> • Date range, if you want a report that includes data through a specific time period. Specify the start and end dates (in MM/DD/YY format) and times (in HH:MM AM/PM 12-Hour format) in the Start Time and Stop Time fields that appear. The maximum time period that you can select is 7 days. • Demand Range for N days, if you want a report that includes data for a specific number of days. Specify the number of days for which demand traffic information must be collected. Range: 1 through 60 days. • Demand Range for last 24 hours, if you want a report that includes data for the past 24 hours.
Report Options > Aggregation Statistic	<p>The traffic is loaded as demand with a configurable number of statistical periods.</p> <p>From the list, select an aggregation statistic option:</p> <ul style="list-style-type: none"> • 90th, 95th, 99th Percentile (X percentile)—Aggregation is based on the selected percentile. The 'X' percentile is the value at which 'X' percent of all the samples taken in the previous sampling period lie at or below the calculated value. • Average—For each interval, the samples within that interval are averaged. If there are N samples for a particular interval, the result is the sum of all the sample values divided by N. • Max—For each interval, the maximum of the sample values within that interval is used. • Min—For each interval, the minimum of the sample values within that interval is used.

Table 178: Fields on the Add Task Wizard (Demand Reports) (Continued)

Field	Description
Report Options > Aggregation Interval	<p>From the list, select the interval for aggregating the collected data:</p> <ul style="list-style-type: none"> • Full Range—The entire range that you specify in the Date Range field is considered as one aggregation interval and one aggregated data point is displayed for the entire range. • Daily— Each day is considered as one aggregation interval and one aggregated data point is displayed per day. • Hourly—Each hour is considered as one aggregation interval and one aggregated data point is displayed per hour.
<i>Schedule</i>	
Startup Options	<p>Schedule—Select one of the following options to schedule the task:</p> <ul style="list-style-type: none"> • Activate Now: The task starts at the current date and time. • Activate Later: The task starts at a later date and time. <p>In the field that appears, click the calendar icon to select the date and time at which you want the task to start. Then, click OK.</p> <p>Instead of using the Activate Now option, you can also click Now at the bottom of the calendar for the task to start at the current date and time.</p>
Recurrence Options	<ul style="list-style-type: none"> • Repeats—Specify the frequency (in Minutes, Hours, Days, Weekly, Monthly, Yearly, or Never) at which the task recurs. • Every—Specify the periodicity of the recurrence. • Ends—Select one of the following options: <ul style="list-style-type: none"> • Never—To configure the task to recur at the specified interval. • On—To select a date and time at which you want the task to end. In the field that appears, click the calendar icon and select the date and time. Then, click OK. • Now—To end a task at the current date and time, click Now at the bottom of the calendar.

RELATED DOCUMENTATION

[Demand Reports Overview | 947](#)

[View Demand Reports | 948](#)

Add a Network Archive Task

From the Task Scheduler page, you can launch collection tasks that create a network model in a database, for use in Paragon Planner. You also have the option to archive the network model.

Tunnel design attributes that are configured in the Tunnel tab of the GUI are inherited by Paragon Planner, even though the attributes are never pushed to the router. When you run the Network Archive task, the tunnel information in Paragon Planner (which came from the router) is merged with the tunnel information in the Path Computation Element (PCE), which includes design attributes that are not pushed to the router. The merged version is then available in Paragon Planner.

The following design attributes that are configured in the Provision LSP pages in the GUI are inherited by Paragon Planner through the network archive task:

- Properties tab: Routing Method
- Advanced tab: Symmetric Pair Group, Diversity Group, and Diversity Level
- Constraints tab: Max Delay, Max Hop, and Max Cost
- Scheduling tab: All scheduling information

To add a network archive task:

1. Select [Administration > Task Scheduler](#).

The Task Scheduler page appears.

2. Click the + (Add) icon.

The Add Task page (wizard) appears.

3. Configure the fields in each step of the wizard according to the guidelines provided in [Table 179 on page 991](#).

Click **Next** to go to the next step.

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

4. After you complete each step in the wizard, click **Finish.**

A confirmation message appears on top of the page, indicating that the task is added successfully.

The details of this task are displayed on the Task Scheduler page.

The network archive files are stored in the PostgreSQL database and can be accessed from there through Paragon Planner. See *Network Browser Window* and *Network Browser Recently Opened and Archived Networks* in the *Paragon Planner User Guide*.

Table 179: Fields on the Add Task Wizard (Network Archive)

Field	Description
<i>Add New Task</i>	
Task Name	Specify a unique name for the task. The name can contain only alphanumeric characters and some special characters (greater than (>), less than (<), colon (:), underscore (_), and hyphen (-)).
Task Group	From the list, select Utility Tasks .
Task Type	From the list, select Network Archive .
<i>Network Archive Task</i>	

Table 179: Fields on the Add Task Wizard (Network Archive) (Continued)

Field	Description
Network Archive Options	<ul style="list-style-type: none"> • Archive network data after processing—Click the toggle button to enable (default) or disable archiving the network data after processing: <ul style="list-style-type: none"> • If you enable this toggle button, the network model created by using the network data is available in Paragon Planner under the Archives tab in the Network Browser window. • If you disable this toggle button, the result of the network archive task is reflected in the new spec file for the latest network archive in Paragon Planner, but it is overwritten by the next latest network archive. • Process Equipment CLI—Equipment CLI data is collected in device collection tasks that include the Equipment CLI option. The Process Equipment CLI option in the Network Archive task parses the Equipment CLI data that is collected during device collection. The inventory reports that are generated are available in both Paragon Pathfinder and Paragon Planner. Click the toggle button to enable (default) or disable this option. <p>To view the inventory reports in Paragon Pathfinder, navigate to Reports > Inventory in the Paragon Automation GUI.</p> <p>To view hardware inventory in Paragon Planner, you must run device collection with the Equipment CLI collection option (collects the inventory data) and the Network Archive collection with the Process Equipment CLI option (processes the inventory data).</p> • Daily Timestamp in Report Title—Click the toggle button to enable or disable (default) the display of the date (in MM:DD:YYYY format) on which the inventory report was generated. If you enable this toggle button, the date is displayed in the Search Reports list in the Inventory Reports page (Reports > Inventory). You can use this option to search for inventory reports generated at a particular date. • Keep Reports for N Days—Specify the number of days until which the inventory reports are retained in the database. <p>Range: 1 through 60.</p> <p>After the specified time elapses, the reports are deleted from the database.</p>

Schedule

Table 179: Fields on the Add Task Wizard (Network Archive) (Continued)

Field	Description
Startup Options	<p>Schedule—Select one of the following options to schedule the task:</p> <ul style="list-style-type: none"> • Activate Now: The task starts at the current date and time. • Activate Later: The task starts at a later date and time. <p>In the field that appears, click the calendar icon to select the date and time at which you want the task to start. Then, click OK.</p> <p>Instead of using the Activate Now option, you can also click Now at the bottom of the calendar for the task to start at the current date and time.</p>
Recurrence Options	<ul style="list-style-type: none"> • Repeats—Specify the frequency (in Minutes, Hours, Days, Weekly, Monthly, Yearly, or Never) at which the task recurs. • Every—Specify the periodicity of the recurrence. • Ends—Select one of the following options: <ul style="list-style-type: none"> • Never—To configure the task to recur at the specified interval. • On—To select a date and time at which you want the task to end. In the field that appears, click the calendar icon and select the date and time. Then, click OK. • Now—To end a task at the current date and time, click Now at the bottom of the calendar.

RELATED DOCUMENTATION

[About the Task Scheduler Page | 969](#)

[Edit and Delete Tasks | 999](#)

Add a Network Maintenance Task

You can add a network maintenance task to schedule maintenance events for network elements, so that you can perform updates or other configuration tasks. Currently, you can add maintenance events only for nodes with the overload bit set.

Maintenance events are planned failures at specific future dates and times. During a scheduled maintenance event, the selected elements are considered logically down. The system reroutes the LSPs around those elements during the maintenance period. After the maintenance event is completed, the default behavior is that all LSPs that were affected by the event are reoptimized.

NOTE: The Path Computation Element (PCE) attempts to reoptimize only PCE-initiated and PCC-delegated LSPs (not PCC-controlled LSPs).

To add a network maintenance task:

1. Select *Administration* > *Task Scheduler*.

The Task Scheduler page appears.

2. Click the + (Add) icon.

The Add Task page (wizard) appears.

3. Configure the fields in each step of the wizard according to the guidelines provided in [Table 180 on page 995](#).

Click **Next** to go to the next step.

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

4. After you complete each step in the wizard, click *Finish*.

A confirmation message appears on top of the page, indicating that the task is added successfully.

The details of this task are displayed on the Task Scheduler page.

Every time the task runs, it first checks the Complete condition for the maintenance event created by the task. If all the elements included in the maintenance task satisfy the Complete condition (for example, `overloadBit = false`), it completes the maintenance event. Next, it looks for elements that match the Create condition (for example, `overloadBit = true`). If it finds such elements, it initiates a new maintenance event that includes those elements.

Just as for other maintenance events, the “M” symbol on the topology map indicates the affected nodes. In the Maintenance tab of the network information table, the maintenance event displays the comment **Created by maintenance task** in the Comment column.

NOTE: This type of maintenance event completes when the included nodes no longer have the overload bit set, but the event will not be deleted automatically. You must manually delete the completed event from the Maintenance tab of the network information table.

Table 180: Fields on the Add Task Wizard (Network Maintenance)

Field	Description
<i>Add New Task</i>	
Task Name	Specify a unique name for the task. The name can contain only alphanumeric characters and some special characters (greater than (>), less than (<), colon (:), underscore (_), and hyphen (-)).
Task Group	From the list, select Utility Tasks .
Task Type	From the list, select Network Maintenance .
<i>Network Maintenance Task</i>	
Task Options	<p>Configure the following fields:</p> <ul style="list-style-type: none"> • Event Name Prefix—Specify a prefix that the PCE uses to name the maintenance event that is initiated by the task. The prefix is followed by a timestamp to ensure that the event name is unique. <p>Alternatively, you can select the check box corresponding to the Use Task Name option to use the name of the task as the prefix.</p> <ul style="list-style-type: none"> • No LSP Optimization Upon Completion—Select the check box corresponding to this option if you don't want the PCE to automatically reoptimize LSPs when the event completes.
Event Conditions	<p>Specify what conditions should trigger the initiation and completion of the maintenance event:</p> <ul style="list-style-type: none"> • Node Bit Overload On Create—Click the toggle button to enable or disable (default) the PCE to automatically initiate a maintenance event for all nodes that have the overload bit set. • Node Bit Overload On Completion—Click the toggle button to enable or disable (default) the PCE to automatically stop the maintenance event for all nodes that have the overload bit set. <p>The PCE discovers the overload bit setting by using either the Network Topology Abstractor Daemon (NTAD) or the BGP Monitoring Protocol (BMP).</p>

Table 180: Fields on the Add Task Wizard (Network Maintenance) (Continued)

Field	Description
<i>Schedule</i>	
Startup Options	<p>Schedule—Select one of the following options to schedule the task:</p> <ul style="list-style-type: none"> • Activate Now: The task starts at the current date and time. • Activate Later: The task starts at a later date and time. <p>In the field that appears, click the calendar icon to select the date and time at which you want the task to start. Then, click OK.</p> <p>Instead of using the Activate Now option, you can also click Now at the bottom of the calendar for the task to start at the current date and time.</p>
Recurrence Options	<ul style="list-style-type: none"> • Repeats—Specify the frequency (in Minutes, Hours, or Days, Weekly, Monthly, Yearly, or Never) at which the task recurs. • Every—Specify the periodicity of the recurrence. • Ends—Select one of the following options: <ul style="list-style-type: none"> • Never—To configure the task to recur at the specified interval. • On—To select a date and time at which you want the task to end. In the field that appears, click the calendar icon and select the date and time. Then, click OK. • Now—To end a task at the current date and time, click Now at the bottom of the calendar.

RELATED DOCUMENTATION

[About the Task Scheduler Page | 969](#)

[Edit and Delete Tasks | 999](#)

Add a Network Cleanup Task

You can add a network cleanup task to clean up the network automatically. Automating this process saves time, especially in large networks.

To add a network cleanup task:

1. Select Administration > Task Scheduler.

The Task Scheduler page appears.

2. Click the + (Add) icon.

The Add Task page (wizard) appears.

3. Configure the fields in each step of the wizard according to the guidelines provided in [Table 181 on page 997](#).

Click **Next** to go to the next step.

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

4. After you complete each step in the wizard, click Finish.

A confirmation message appears on top of the page, indicating that the task is added successfully.

The details of this task are displayed on the Task Scheduler page.

You can view the post-cleanup topology on the Topology page (**Network > Topology**). To ensure that you have the latest topology, right-click the blank space in the topology map and select **Reload Network**. The latest topology is displayed in the topology map.

Table 181: Fields on the Add Task Wizard (Network Cleanup)

Field	Description
<i>Add New Task</i>	
Task Name	Specify a unique name for the task. The name can contain only alphanumeric characters and some special characters (greater than (>), less than (<), colon (:), underscore (_), and hyphen (-)).
Task Group	From the list, select Utility Tasks .
Task Type	From the list, select Network Cleanup .

Table 181: Fields on the Add Task Wizard (Network Cleanup) (Continued)

Field	Description
<i>Network Cleanup Task</i>	
Network Cleanup Options	<p>Select one or more of the following options. By default, all the available options are selected except Force remove links with user attributes.</p> <ul style="list-style-type: none"> • Purge down links—If you select this option, the links that are down are removed from the live network and database. • Force remove links with user attributes—If you select this option, the links that have custom attributes defined in the controller are removed from the live network and database. • Purge down nodes—If you select this option, the nodes that are down are removed from the live network and database. • Generate purge report—If you select this option, a report is generated every time the task is executed. The report indicates the actions taken as a result of the cleanup. Purge reports, identified with a timestamp, are stored in <code>/opt/northstar/data/.network_plan/Report/purge_reports/</code> • Add notifications to timeline—If you select this option, you can see notifications relevant to the execution of the task in the Timeline page. To get there, navigate to the Topology page (Network > Topology) and right-click on the blank space on the topology map. From the list that appears, select Timeline.
<i>Schedule</i>	
Startup Options	<p>Schedule—Select one of the following options to schedule the task:</p> <ul style="list-style-type: none"> • Activate Now: The task starts at the current date and time. • Activate Later: The task starts at a later date and time. <p>In the field that appears, click the calendar icon to select the date and time at which you want the task to start. Then, click OK.</p> <p>Instead of using the Activate Now option, you can also click Now at the bottom of the calendar for the task to start at the current date and time.</p>

Table 181: Fields on the Add Task Wizard (Network Cleanup) (Continued)

Field	Description
Recurrence Options	<ul style="list-style-type: none"> • Repeats—Specify the frequency (in Minutes, Hours, Days, Weekly, Monthly, Yearly, or Never) at which the task recurs. • Every—Specify the periodicity of the recurrence. • Ends—Select one of the following options: <ul style="list-style-type: none"> • Never—To configure the task to recur at the specified interval. • On—To select a date and time at which you want the task to end. In the field that appears, click the calendar icon and select the date and time. Then, click OK. • Now—To end a task at the current date and time, click Now at the bottom of the calendar.

RELATED DOCUMENTATION

[About the Task Scheduler Page | 969](#)

Edit and Delete Tasks

IN THIS SECTION

- [Edit Tasks | 1000](#)
- [Delete Tasks | 1000](#)

You can modify (edit) the parameters configured for existing tasks and delete tasks that are no longer needed.

NOTE: You can only edit and delete user-created tasks.

Edit Tasks

To edit a task:

1. Select **Administration > Task Scheduler**.
The Task Scheduler page appears.
2. Select the task that you want to modify and click the **Edit** (Pencil) icon.
The Edit Task page appears.
3. Modify the parameters as required.

NOTE: You cannot modify the name, task group, and task type for the tasks.

4. Click **Finish** to save your changes.
The modifications are saved. You are returned to the Task Scheduler page, where a confirmation message appears.

Delete Tasks

To delete one or more tasks:

1. Select **Administration > Task Scheduler**.
The Task Scheduler page appears.
2. Select one or more tasks that you want to delete, and click the **Delete** (trash can) icon.
An alert message appears asking you to confirm the deletion.
3. Click **Yes**.
The selected tasks are deleted. You are returned to the Task Scheduler page, where a confirmation message appears.

RELATED DOCUMENTATION

| [About the Task Scheduler Page](#) | 969

Manage Security Settings

IN THIS CHAPTER

- [About the Security Settings Page | 1001](#)
- [Configure Security Profiles for SSL Authentication | 1003](#)

About the Security Settings Page

IN THIS SECTION

- [Tasks You Can Perform | 1002](#)
- [Field Descriptions | 1002](#)

To access this page from the Paragon Automation graphical user interface (GUI), click **Administration > Security**.

Each time you navigate to the **Security** page, you first see the **CA Profiles** page. Click the tabs on the right to toggle between **Local Certificate** and **CA Profile** tabbed pages.

In Paragon Automation, you can create CA Profiles and Local Certificate profiles for secure authentication in AMQP-based notifications. Paragon Insights supports various authentication methods to provide secure data connection for Paragon Insights devices. [Table 182 on page 1002](#) provides an overview of supported authentication methods and security parameters.

Table 182: Paragon Insights Authentication Methods

Authentication Method	Description	Required Paragon Insights Security Parameters
Mutual SSL	Client authenticates itself with the server and the server authenticates itself with the client.	<ul style="list-style-type: none"> Local certificates (includes the client certificate and client key) CA certificate Server common name
Server-side SSL	Server authenticates itself with the client.	<ul style="list-style-type: none"> CA certificate Server common name
Password	Authenticates users with a password.	<ul style="list-style-type: none"> Username Password

Tasks You Can Perform

You can perform the following tasks from this page:

- Create and manage CA Profiles.
- Create and manage Local Certificates.

Field Descriptions

[Table 183 on page 1002](#) describes the fields on the CA Profiles page.

[Table 184 on page 1003](#) describes the fields on the Local Certificates page.

Table 183: Fields on the CA Profiles Page

Field	Description
Name	View the name of the CA Profile.

Table 183: Fields on the CA Profiles Page (Continued)

Field	Description
Certificate	View the uploaded certificate filename.

Table 184: Fields on the Local Certificate Page

Field	Description
Name	View the name of the Local Certificate profile.
Certificate	View the uploaded client certificate filename.
Key	View the uploaded client key filename.

RELATED DOCUMENTATION

| [Configure Security Profiles for SSL Authentication](#) | 1003

Configure Security Profiles for SSL Authentication

You can configure security profiles for (SSL) authentication from the Paragon Automation graphical user interface (GUI).

For more information on security settings, see "[About the Security Settings Page](#)" on page 1001.

To configure security profiles for SSL authentication:

1. Click **Administration** > **Security**.
The **Security Settings** page is displayed.
2. To add a CA profile:
 - a. Click the **CA Profiles** tab.
The **CA Profiles** page is displayed.

- b. Click (+) icon to add a CA profile.
The **Add CA Profiles** page is displayed.

- c. Enter the following information:

Name	Enter a name for the CA profile.
Upload Certificate	Click Choose file and navigate to the location of the file that you want to upload. Select the CA certificate file and then click Open . The supported file extension is .crt .

To add a local certificate profile:

- a. Click the **Local Certificates** tab.
The **Local Certificates** page is displayed.
- b. Click (+) icon to add a new local certificate profile.
The **Add Local Certificate Profile** page is displayed.

- c. Enter the following information:

Name	Enter a name for the local certificate profile.
Upload Certificate	Click Choose File and navigate to the location of the client certificate file. Select the file and then click Open . The supported file extension is .crt .
Upload Key	Click Choose File and navigate to the location of the client key file. Select the file and then click Open . The supported file extension is .key .

- 3. Click **Save** to only save the configuration.
Click **Save and Deploy** to save and deploy the configuration immediately.

RELATED DOCUMENTATION

| [About the Security Settings Page](#) | **1001**

License Management

IN THIS CHAPTER

- [Paragon Insights Licensing Overview | 1005](#)
- [About the License Management Page | 1006](#)
- [View, Add, or Delete Licenses | 1009](#)

Paragon Insights Licensing Overview

Juniper Networks introduced the Juniper Flex Software Subscription Licensing model to provide an efficient way for you to manage licenses for hardware and software features. Paragon Insights uses the Juniper Flex Software Subscription Licensing model. For more information, see the [Paragon Insights Licensing](#) topic in the Licensing Guide.

You need a license to activate the Graphical User Interface (GUI). When you log in to the Paragon Automation GUI for the first time, the **Dashboard** page appears. Navigate to **Administration > License Management** to add a license. After you successfully add a license for a component (Paragon Insights, Paragon Pathfinder, or Paragon Planner), you can see the related GUI pages. The availability of features in Paragon Insights, Paragon Pathfinder, and Paragon Planner are based on the license you have purchased.

The Paragon Automation icon in the top-left corner of the GUI is updated depending on the license that you add. When you log in to the Paragon Automation GUI for the first time, the Paragon Automation icon appears without the names of the three components below it. The GUI displays the name of a component only after you add a license for that component. For example, after you add a Paragon Insights license, the name **Insights** appears below the Paragon Automation icon. After you add a license for Paragon Pathfinder, the names **Pathfinder** and **Planner** are also displayed.

RELATED DOCUMENTATION

[About the License Management Page | 1006](#)

[View, Add, or Delete Licenses | 1009](#)

About the License Management Page

IN THIS SECTION

- [Tasks You Can Perform | 1006](#)
- [Field Descriptions | 1006](#)

To access this page from the Paragon Automation GUI, click **Administration > License Management**.

Juniper Networks introduced the Juniper Flex Software Subscription Licensing model to provide an efficient way for you to manage licenses for hardware and software features.

The **License Management Page** provides a complete view of the available licensing features, and details of all licenses added.

You need a license to activate the GUI. When you log in to the Paragon Automation GUI for the first time, the **Dashboard** page appears. Navigate to **Administration > License Management** to add a license. After you successfully add a license for a component (Paragon Insights, Paragon Pathfinder, or Paragon Planner), you can see the related GUI pages. The availability of features in Paragon Insights, Paragon Pathfinder, and Paragon Planner are based on the license you have purchased.

Tasks You Can Perform

You can perform the following tasks from this page:

- Add a license.
- Delete a license.
- View licensing features.
- View status and details of a license.

Field Descriptions

[Table 185 on page 1007](#) describes the fields of the Feature Summary section.

[Table 186 on page 1007](#) describes the fields of the Added License section.

Table 185: View Details of Feature Summary Section

Attribute	Description
Feature	View the names of the licenses.
Description	View a brief description of the license.
License Limit	View the number of valid licenses successfully added and available for use. NOTE: Irrespective of the number of PIN-Advanced, and PIN-Standard platform licenses that you add, the maximum license limit is 1 . However, with PIN-Devices licenses, the limit changes with the number of licenses that you add.
Usage Count	View the number of available licenses that are currently in use.
Valid Until	View the date and time of when the license expires.
Compliance	View color definitions that determine license compliance. Green Feature licenses are in compliance with Juniper's End User License Agreement. Yellow Device feature licenses are $\geq 90\%$ of the limit. You are getting close to running out of licenses. This status is applies only to device feature licenses. Red Feature licenses are not in compliance with Juniper's End User License Agreement. Click the red dot to view details about the compliance issue.

Table 186: View Details of Added License Section

Attribute	Description
License ID	View the license identification number generated through the Juniper Agile Licensing Portal.

Table 186: View Details of Added License Section (*Continued*)

Attribute	Description
Customer ID	View the customer identification. NOTE: The customer ID might not be displayed after you add a license. This is because the customer ID is not embedded in the new license key format.
Order Type	View the order type (commercial, demo, education, emergency, lab, and unknown).
Validity Type	View the validity type of a license (date-based or permanent).
Start Date	View the start date of the license.
End Date	View the end date of the license.
State	Displays the state of the license.
Feature ID	View the feature license identification number.
Feature Name	View the feature license name.
Feature Description	View a brief description of the feature license.
License Count	View the entitled license count for this feature license.

NOTE: In Paragon Automation Release 22.1, **SKU Name** is not available in the **Added License** section of the **License Management** page. This is because the new license key format does not support **SKU Name**.

RELATED DOCUMENTATION

[Paragon Insights Licensing Overview | 1005](#)

View, Add, or Delete Licenses

IN THIS SECTION

- [Add a License | 1010](#)
- [Delete a License | 1010](#)
- [View Licensing Features | 1011](#)
- [View Status and Details of a License | 1012](#)

When you navigate to the **License Management** page, a warning message asking you to add a tier-based license or upgrade to the new license key format might be displayed.

- If no licenses are added, you must obtain a license key from the [Juniper Agile Licensing Portal](#) and add it to the License Management page.
- If licenses added but the license key format is not compatible with the Paragon Automation software version that you are currently using, you must generate a new license key.

To generate a new license key in the compatible format:

1. Login to [Juniper Agile Licensing Portal](#).
2. Revoke the existing license.

A new license key is generated.

NOTE:

Issue DateLicense Key

3. Activate the new license key that is compatible with the software version you're running. Select the software version during the activation process.
For more information, see the FAQ section in the [Juniper Agile Licensing Portal](#).
4. Add the new license key to the **License Management** page.
See "[Add a License](#)" on page 1010.

After you have obtained a license through the [Juniper Agile Licensing Portal](#), you can manage your licenses from the Paragon Automation GUI .

The **Administration > License Management** page enables you to:

- Add a license.
- Delete a license.
- View licensing features.
- View status and details of a license.

Add a License

To add a license:

1. Navigate to **Administration > License Management** page.
The **License Management** page appears.
2. Click (+) icon in the **Licenses Added** section to add a new license.
The **Add License** pop-up appears.
3. Click **Browse** and navigate to the location of the license file that you want to add.
Select the file and then click **Open**.

The license file is added to the **Add License** page.
4. Click **Ok** to confirm selection.
The license you added is listed in the **Licenses Added** section of the **License Management** page.

Delete a License

To delete an existing license:

1. Navigate to **Administration > License Management** page.
The **License Management** page appears.
2. From the **Licenses Added** section, click the option button (available in beginning the row) of the license you want to delete, and then click **Delete**.

The **Confirm Delete** page appears.
3. Click **Yes** to delete the license.

View Licensing Features

To view licensing features, navigate to **Administration > License Management** page.

The Features Summary section, as given in [Table 187 on page 1011](#) , provides information on the feature licensing attributes.

Table 187: Feature Licensing Attributes

Attribute	Description
Feature	View the name of the licenses.
Description	View a brief description of the feature license.
License Limit	View the number of valid licenses successfully added and available for use. NOTE: Irrespective of the number of PIN-Advanced, and PIN-Standard platform licenses that you add, the maximum license limit is 1 . However, with PIN-Devices licenses, the limit changes with the number of licenses that you add.
Usage Count	View the number of available licenses that are currently in use in this instance of Paragon Automation.
Valid Until	View the date and time when the license expires.
Compliance	Color definitions (dot indicator) that determine license compliance: Green Feature licenses are in compliance with Juniper's End User License Agreement. Yellow Device feature licenses are $\geq 90\%$ of the limit. You are getting close to running out of licenses. This status applies only to device feature licenses. Red Feature licenses are not in compliance with Juniper's End User License Agreement. Click the red dot to view details about the compliance issue.

View Status and Details of a License

To view status and details of a license, navigate to **Administration > License Management** page. The **Added Licenses** section provides information about the added licenses. You can also click the option button (available in beginning the row) of the license you want to view, and then click the **More > Details** to view more information about the license.

[Table 188 on page 1012](#) provides information on the feature licensing attributes.

Table 188: View Details of Licenses Added

Attribute	Description
License ID	Identification number for the license generated through the Juniper Agile Licensing Portal.
Customer ID	Identification for the customer. NOTE: The customer ID might not be displayed after you add a license. This is because the customer ID is not embedded in the new license key format.
Order Type	View the order type (commercial, demo, education, emergency, lab, and unknown).
Validity Type	View the validity type of a license (date-based or permanent).
Start Date	View the start date of the license.
End Date	View the end date of the license.
State	Displays the state of the license.
Feature ID	View the identification number for the feature license.
Feature Name	View the feature license name.
Feature Description	View a brief description of the Paragon Insights feature license.

Table 188: View Details of Licenses Added (Continued)

Attribute	Description
License Count	View the entitled license count for this feature license.

NOTE: In Paragon Automation Release 22.1, **SKU Name** is not available in the **Added License** section of the **License Management** page. This is because the new license key format does not support **SKU Name**.

RELATED DOCUMENTATION

[About the License Management Page | 1006](#)

[Paragon Insights Licensing Overview | 1005](#)